

Clustering and Classification via Cluster-Weighted Factor Analyzers

Sanjeena Subedi ^a, Antonio Punzo ^b, Salvatore Ingrassia ^{b,*},
Paul D. McNicholas ^a

^a*Department of Mathematics and Statistics, University of Guelph
Ontario, Canada, N1G 2W1.*

^b*Department of Economics and Business, University of Catania
Corso Italia 55, 95129 Catania, Italy.*

Abstract

In model-based clustering and classification, the cluster-weighted model constitutes a convenient approach when the random vector of interest constitutes a response variable Y and a set p of explanatory variables \mathbf{X} . However, its applicability may be limited when p is high. To overcome this problem, this paper assumes a latent factor structure for \mathbf{X} in each mixture component. This leads to the cluster-weighted factor analyzers (CWFA) model. By imposing constraints on the variance of Y and the covariance matrix of \mathbf{X} , a novel family of sixteen CWFA models is introduced for model-based clustering and classification. The alternating expectation-conditional maximization algorithm, for maximum likelihood estimation of the parameters of all the models in the family, is described; to initialize the algorithm, a 5-step hierarchical procedure is proposed, which uses the nested structures of the models within the family and thus guarantees the natural ranking among the sixteen likelihoods. Artificial and real data show that these models have very good clustering and classification performance and that the algorithm is able to recover the parameters very well.

Key words: Cluster-weighted models; factor analysis; mixture models; parsimonious models.

* Corresponding author

Email addresses: ssubedi@uoguelph.ca (Sanjeena Subedi),
antonio.punzo@unict.it (Antonio Punzo), s.ingrassia@unict.it (Salvatore Ingrassia), paul.mcnicholas@uoguelph.ca (Paul D. McNicholas).

1 Introduction

In direct applications of finite mixture models, each of the G mixture components is taken to represent a sub-group (or cluster) within the data (see Titterington et al., 1985, pp. 2–3). The terms ‘model-based clustering’ and ‘model-based classification’ have been used to describe the adoption of mixture models or, more often, a family of mixture models for clustering and classification, respectively. In the 1990’s, three model-based clustering papers (Banfield and Raftery, 1993, Celeux and Govaert, 1995, and Ghahramani and Hinton, 1987) effectively set the scene for the push towards the finite mixture model-based approaches that followed. Overviews of mixture models and their applications are given in Everitt and Hand (1981), Titterington et al. (1985), McLachlan and Peel (2000), and Frühwirth-Schnatter (2006).

Consider a random vector $(\mathbf{X}', Y)'$, defined from Ω to $\mathbb{R}^p \times \mathbb{R}$, where a latent group-structure as well as a linear dependence of Y on \mathbf{x} in each group are assumed. Under these assumptions, the linear cluster-weighted model (CWM; introduced in Gershensfeld, 1997) is an ideal choice within the mixture modelling framework. It factorizes the joint density of $(\mathbf{X}', Y)'$, in each mixture-component, into the product of the conditional density of $Y|\mathbf{x}$ and the marginal density of \mathbf{X} . In this manner, the model takes into account the potential of finite mixtures of regressions (see Frühwirth-Schnatter, 2006, Chapter 8) in modelling the conditional density of $Y|\mathbf{x}$, and the potential of finite mixtures of Gaussian distributions (see Titterington et al., 1985 and McLachlan and Peel, 2000) in modelling the joint density of $(\mathbf{X}', Y)'$ and the marginal density of \mathbf{X} .

Recent literature on model-based clustering and classification through the CWM can be summarized as follows. Ingrassia, Minotti, and Vittadini (2012) study the relationships between the linear Gaussian CWM and some well-known mixture-based approaches, moreover they consider the t -distribution as a robust alternative to Gaussian assumptions. By using this model as a building block, Ingrassia, Minotti, and Punzo (2012) introduce a family of parsimonious linear t -CWMs for model-based clustering. Finally, under Gaussian assumptions for both mixture-component densities, Punzo (2012) introduces the polynomial CWM as a flexible tool for clustering and classification purposes.

However, the applicability of linear Gaussian CWMs in high dimensional \mathbf{X} -spaces still remains a challenge. The number of parameters for this model is $(G - 1) + G(p + 2) + G[p + p(p + 1)/2]$, of which $Gp(p + 1)/2$ are used for the group covariance matrices Σ_g of \mathbf{X} alone, $g = 1, \dots, G$, and this increases quadratically with p . To overcome this issue, we assume a latent Gaussian factor structure for \mathbf{X} , in each mixture-component, which leads to the *Fac-*

tor Regression Model (FRM) of Y on \mathbf{x} (see West, 2003, Wang et al., 2007, and Carvalho et al., 2008). The FRM assumes $\Sigma_g = \Lambda_g \Lambda_g' + \Psi_g$, where the loading matrix is a $p \times q$ matrix of parameters typically with $q \ll p$ and the noise matrix Ψ_g is a diagonal matrix. The adoption of this group covariance structure in the linear Gaussian CWM framework leads to the linear Gaussian cluster-weighted factor analyzers model (CWFA), which is characterized by $G[pq - q(q-1)/2] + Gp$ parameters for the group covariance matrices. The CWFA model follows the principle of the general form of mixtures of factor analyzers regarding \mathbf{X} ; mixtures of factor analyzers were developed by McLachlan and Peel (2000, Chapter 8) and McLachlan et al. (2003) on the basis of the original work of Ghahramani and Hinton (1987). Furthermore, starting from the works of McNicholas and Murphy (2008), McNicholas (2010), Ingrassia, Minotti, and Vittadini (2012) and Ingrassia, Minotti, and Punzo (2012), a novel family of sixteen mixture models — obtained as special cases of the linear Gaussian CWFA by conveniently constraining the component variances of Y and \mathbf{X} — is introduced to facilitate parsimonious model-based clustering and classification in the defined paradigm.

The paper is organized as follows. Section 2 recalls the linear Gaussian CWM and the FRM (with details given in Appendix A); they are the basic models to define the linear Gaussian CWFA models introduced in Section 3. Model fitting with the alternating expectation-conditional maximization (AECM) algorithm is presented in Section 4, with details given in Appendix B.2. Section 5 addresses computational details on some aspects of the AECM algorithm and discusses model selection and evaluation. Artificial and real data are considered in Section 6, and the paper concludes with discussion and suggestions for further work in Section 7.

2 Model

This section provides a step-by-step introduction to the model we introduce in the next section.

2.1 The linear Gaussian cluster-weighted model

Let $p(\mathbf{x}, y)$ be the joint density of $(\mathbf{X}', Y)'$. Suppose that Ω can be partitioned into G groups, say $\Omega_1, \dots, \Omega_G$. The CWM defines the joint density as

$$p(\mathbf{x}, y; \boldsymbol{\theta}) = \sum_{g=1}^G \pi_g p(y|\mathbf{x}, \Omega_g) p(\mathbf{x}|\Omega_g), \quad (1)$$

where $p(y|\mathbf{x}, \Omega_g)$ is the conditional density of the response variable Y given \mathbf{x} and Ω_g , $p(\mathbf{x}|\Omega_g)$ is the marginal density of \mathbf{x} given Ω_g , $\pi_g = p(\Omega_g)$ is the weight of Ω_g in the mixture (defined so that $\pi_g > 0$ and $\sum_g \pi_g = 1$), $g = 1, \dots, G$, and $\boldsymbol{\theta}$ contains all of the parameters in the mixture.

The component densities $p(\mathbf{x}|\Omega_g)$ and $p(y|\mathbf{x}, \Omega_g)$ are usually assumed to be (multivariate) Gaussian (see, e.g., Ingrassia, Minotti, and Vittadini, 2012 and Punzo, 2012), the former with mean vector $\boldsymbol{\mu}_g$ and covariance matrix $\boldsymbol{\Sigma}_g$ and the latter with linear conditional mean $\mu(\mathbf{x}, \boldsymbol{\beta}_g) = \beta_{0g} + \boldsymbol{\beta}'_{1g}\mathbf{x}$ and conditional variance σ_g^2 , where $\boldsymbol{\beta}_g = (\beta_{0g}, \boldsymbol{\beta}'_{1g})'$, $\beta_{0g} \in \mathbb{R}$, and $\boldsymbol{\beta}_{1g} \in \mathbb{R}^p$. In other words, conditional on \mathbf{x} and Ω_g , the linear model $Y|\mathbf{x} = \mu(\mathbf{x}, \boldsymbol{\beta}_g) + \varepsilon_g$ holds. Thus, the general CWM in (1) becomes the linear Gaussian CWM

$$p(\mathbf{x}, y; \boldsymbol{\theta}) = \sum_{g=1}^G \pi_g \phi(y|\mathbf{x}; \mu(\mathbf{x}, \boldsymbol{\beta}_g), \sigma_g^2) \phi(\mathbf{x}; \boldsymbol{\mu}_g, \boldsymbol{\Sigma}_g), \quad (2)$$

where

$$\begin{aligned} \phi(y|\mathbf{x}; \mu(\mathbf{x}, \boldsymbol{\beta}_g), \sigma_g^2) &= \frac{1}{\sqrt{2\pi\sigma_g^2}} \exp \left\{ -\frac{[y - \mu(\mathbf{x}, \boldsymbol{\beta}_g)]^2}{2\sigma_g^2} \right\}, \\ \phi(\mathbf{x}; \boldsymbol{\mu}_g, \boldsymbol{\Sigma}_g) &= \frac{1}{(2\pi)^{p/2} |\boldsymbol{\Sigma}_g|^{p/2}} \exp \left\{ -\frac{1}{2} (\mathbf{x} - \boldsymbol{\mu}_g)' \boldsymbol{\Sigma}_g^{-1} (\mathbf{x} - \boldsymbol{\mu}_g) \right\}. \end{aligned}$$

2.2 The factor regression model

The factor analysis model (Spearman, 1904 and Bartlett, 1953), for the p -dimensional variable \mathbf{X} , postulates that

$$\mathbf{X} = \boldsymbol{\mu} + \boldsymbol{\Lambda}\mathbf{U} + \mathbf{e}, \quad (3)$$

where $\mathbf{U} \sim N_q(\mathbf{0}, \mathbf{I}_q)$ is a q -dimensional ($q \ll p$) vector of latent factors, $\boldsymbol{\Lambda}$ is a $p \times q$ matrix of factor loadings, and $\mathbf{e} \sim N_p(\mathbf{0}, \boldsymbol{\Psi})$, with $\boldsymbol{\Psi} = \text{diag}(\psi_1^2, \dots, \psi_p^2)$, independent of \mathbf{U} . Then $\mathbf{X} \sim N_p(\boldsymbol{\mu}, \boldsymbol{\Lambda}\boldsymbol{\Lambda}' + \boldsymbol{\Psi})$ and, conditional on \mathbf{u} , results in $\mathbf{X}|\mathbf{u} \sim N_p(\boldsymbol{\mu} + \boldsymbol{\Lambda}\mathbf{u}, \boldsymbol{\Psi})$.

Model (3) can be considered similarly to the standard (linear) regression model $Y = \beta_0 + \boldsymbol{\beta}'_1\mathbf{X} + \varepsilon$ leading to the FRM (see West, 2003, Wang et al., 2007, and Carvalho et al., 2008)

$$Y = \beta_0 + \boldsymbol{\beta}'_1(\boldsymbol{\mu} + \boldsymbol{\Lambda}\mathbf{U} + \mathbf{e}) + \varepsilon = (\beta_0 + \boldsymbol{\beta}'_1\boldsymbol{\mu}) + \boldsymbol{\beta}'_1\boldsymbol{\Lambda}\mathbf{U} + (\boldsymbol{\beta}'_1\mathbf{e} + \varepsilon),$$

where ε is assumed to be independent of \mathbf{U} and \mathbf{e} . The mean and variance of

Y are given by

$$\begin{aligned}\mathbb{E}(Y) &= \beta_0 + \beta_1' \boldsymbol{\mu} \\ \text{Var}(Y) &= \text{Var}(\beta_1' \boldsymbol{\Lambda} \mathbf{U}) + \text{Var}(\beta_1' \mathbf{e}) + \text{Var}(\varepsilon) \\ &= \beta_1' \boldsymbol{\Lambda} \boldsymbol{\Lambda}' \beta_1 + \beta_1' \boldsymbol{\Psi} \beta_1 + \sigma^2 = \beta_1' (\boldsymbol{\Lambda} \boldsymbol{\Lambda}' + \boldsymbol{\Psi}) \beta_1 + \sigma^2,\end{aligned}$$

respectively, and so $Y \sim N(\beta_0 + \beta_1' \boldsymbol{\mu}, \beta_1' (\boldsymbol{\Lambda} \boldsymbol{\Lambda}' + \boldsymbol{\Psi}) \beta_1 + \sigma^2)$.

Consider the triplet $(Y, \mathbf{X}', \mathbf{U}')'$. Its mean is given by

$$\mathbb{E} \begin{bmatrix} Y \\ \mathbf{X} \\ \mathbf{U} \end{bmatrix} = \begin{bmatrix} \beta_0 + \beta_1' \boldsymbol{\mu} \\ \boldsymbol{\mu} \\ \mathbf{0} \end{bmatrix},$$

and because $\text{Cov}(\mathbf{X}, Y) = (\boldsymbol{\Lambda} \boldsymbol{\Lambda}' + \boldsymbol{\Psi}) \beta_1$ and $\text{Cov}(\mathbf{U}, Y) = \boldsymbol{\Lambda}' \beta_1$, it results

$$\text{Cov} \begin{bmatrix} Y \\ \mathbf{X} \\ \mathbf{U} \end{bmatrix} = \begin{bmatrix} \beta_1' \boldsymbol{\Sigma} \beta_1 + \sigma^2 & \beta_1' \boldsymbol{\Sigma} & \beta_1' \boldsymbol{\Lambda} \\ \boldsymbol{\Sigma} \beta_1 & \boldsymbol{\Sigma} & \boldsymbol{\Lambda} \\ \boldsymbol{\Lambda}' \beta_1 & \boldsymbol{\Lambda}' & \mathbf{I}_q \end{bmatrix},$$

where $\boldsymbol{\Sigma} = \boldsymbol{\Lambda} \boldsymbol{\Lambda}' + \boldsymbol{\Psi}$. Now, we can write the joint density of $(Y, \mathbf{X}', \mathbf{U}')'$ as

$$p(y, \mathbf{x}, \mathbf{u}) = \phi(y|\mathbf{x}, \mathbf{u}) \phi(\mathbf{x}|\mathbf{u}) \phi(\mathbf{u}). \quad (4)$$

Here, the distribution and related parameters for both $\mathbf{X}|\mathbf{u}$ and \mathbf{U} are known. Thus, we need only to analyze the distribution of $Y|\mathbf{x}, \mathbf{u}$. Importantly, $\mathbb{E}(Y|\mathbf{x}, \mathbf{u}) = \mathbb{E}(Y|\mathbf{x})$ and $\text{Var}(Y|\mathbf{x}, \mathbf{u}) = \text{Var}(Y|\mathbf{x})$, and so $Y|\mathbf{x}, \mathbf{u} \sim N(\beta_0 + \beta_1' \mathbf{x}, \sigma^2)$; mathematical details are given in Appendix A. This implies that $\phi(y|\mathbf{x}, \mathbf{u}) = \phi(y|\mathbf{x})$ and, therefore, Y is conditionally independent of \mathbf{U} given $\mathbf{X} = \mathbf{x}$, so that (4) becomes

$$p(y, \mathbf{x}, \mathbf{u}) = \phi(y|\mathbf{x}) \phi(\mathbf{x}|\mathbf{u}) \phi(\mathbf{u}). \quad (5)$$

Similarly, $\mathbf{U}|y, \mathbf{x} \sim N(\boldsymbol{\gamma}(\mathbf{x} - \boldsymbol{\mu}), \mathbf{I}_q - \boldsymbol{\gamma} \boldsymbol{\Lambda})$, where $\boldsymbol{\gamma} = \boldsymbol{\Lambda}' (\boldsymbol{\Lambda} \boldsymbol{\Lambda}' + \boldsymbol{\Psi})^{-1}$, and thus \mathbf{U} is conditionally independent on Y given $\mathbf{X} = \mathbf{x}$. Therefore,

$$\begin{aligned}\mathbb{E}[\mathbf{U}|\mathbf{x}; \boldsymbol{\mu}, \boldsymbol{\Lambda}, \boldsymbol{\Psi}] &= \boldsymbol{\gamma}(\mathbf{x} - \boldsymbol{\mu}), \quad \text{and} \\ \mathbb{E}[\mathbf{U} \mathbf{U}'|\mathbf{x}; \boldsymbol{\mu}, \boldsymbol{\Lambda}, \boldsymbol{\Psi}] &= \mathbf{I}_q - \boldsymbol{\gamma} \boldsymbol{\Lambda} + \boldsymbol{\gamma}(\mathbf{x} - \boldsymbol{\mu})(\mathbf{x} - \boldsymbol{\mu})' \boldsymbol{\gamma}'.\end{aligned}$$

3 The modelling framework

3.1 The general model

Assume that for each Ω_g , $g = 1, \dots, G$, the pair $(\mathbf{X}', Y)'$ satisfies a FRM, that is

$$Y = \beta_{0g} + \beta'_{1g}\mathbf{X} + \varepsilon_g \quad \text{with} \quad \mathbf{X} = \boldsymbol{\mu}_g + \boldsymbol{\Lambda}_g\mathbf{U}_g + \mathbf{e}_g, \quad (6)$$

where $\boldsymbol{\Lambda}_g$ is a $p \times q$ matrix of factor loadings, $\mathbf{U}_g \sim N_q(\mathbf{0}, \mathbf{I}_q)$ is the vector of factors, $\mathbf{e}_g \sim N_p(\mathbf{0}, \boldsymbol{\Psi}_g)$ are the errors, $\boldsymbol{\Psi}_g = \text{diag}(\psi_{1g}, \dots, \psi_{pg})$, and $\varepsilon_g \sim N(0, \sigma_g^2)$. Then the linear Gaussian CWM in (2) can be extended in order to include the underlying factor structure (6) for the \mathbf{X} variable. In particular, by recalling that Y is conditionally independent of \mathbf{U} given $\mathbf{X} = \mathbf{x}$ in the generic Ω_g , we get

$$p(\mathbf{x}, y; \boldsymbol{\theta}) = \sum_{g=1}^G \pi_g \phi(y|\mathbf{x}; \mu(\mathbf{x}; \boldsymbol{\beta}_g), \sigma_g^2) \phi(\mathbf{x}; \boldsymbol{\mu}_g, \boldsymbol{\Lambda}_g\boldsymbol{\Lambda}_g' + \boldsymbol{\Psi}_g), \quad (7)$$

where $\boldsymbol{\theta} = \{\pi_g, \boldsymbol{\beta}_g, \sigma_g^2, \boldsymbol{\mu}_g, \boldsymbol{\Lambda}_g, \boldsymbol{\Psi}_g; g = 1, \dots, G\}$. Model (7) is the linear Gaussian CWFA, which we shall refer to as the CWFA model herein.

3.2 Parsimonious versions of the model

In this section, we extend the linear Gaussian CWFA by allowing constraints across groups on σ_g^2 , $\boldsymbol{\Lambda}_g$, and $\boldsymbol{\Psi}_g$, and on whether or not $\boldsymbol{\Psi}_g = \psi_g\mathbf{I}_p$ (isotropic assumption). The full range of possible constraints provides a family of sixteen different parsimonious CWFAs, which are given in Table 1.

Here, models are identified by a sequence of four letters. The letters refer to whether or not the constraints $\sigma_g^2 = \sigma^2$, $\boldsymbol{\Lambda}_g = \boldsymbol{\Lambda}$, $\boldsymbol{\Psi}_g = \boldsymbol{\Psi}$, and $\boldsymbol{\Psi}_g = \psi_g\mathbf{I}_p$, respectively, are imposed. The constraints on the group covariances of \mathbf{X} are in the spirit of McNicholas and Murphy (2008), while that on the group variances of Y are borrowed from Ingrassia, Minotti, and Punzo (2012). Each letter can be either C, if the corresponding constraint is applied, or U if the particular constraint is not applied. For example, model CUUC assumes equal Y variances between groups, unequal loading matrices, and unequal, but isotropic, noise.

3.3 Model-based classification

Suppose that m of the n observations in \mathcal{S} are labeled. Within the model-based classification framework, we use all of the n observations to estimate

Table 1

Parsimonious covariance structures derived from the CWFA model.

Model ID	Y Variance	Loading Matrix	Error Variance	Isotropic	Covariance parameters
UUUU	unconstrained	unconstrained	unconstrained	unconstrained	$G + G[pq - q(q-1)/2] + Gp$
UUUC	unconstrained	unconstrained	unconstrained	constrained	$G + G[pq - q(q-1)/2] + G$
UUCU	unconstrained	unconstrained	constrained	unconstrained	$G + G[pq - q(q-1)/2] + p$
UUCU	unconstrained	unconstrained	constrained	constrained	$G + G[pq - q(q-1)/2] + 1$
UCUU	unconstrained	constrained	unconstrained	unconstrained	$G + [pq - q(q-1)/2] + Gp$
UCUC	unconstrained	constrained	unconstrained	constrained	$G + [pq - q(q-1)/2] + G$
UCCU	unconstrained	constrained	constrained	unconstrained	$G + [pq - q(q-1)/2] + p$
UCCC	unconstrained	constrained	constrained	constrained	$G + [pq - q(q-1)/2] + 1$
CUUU	constrained	unconstrained	unconstrained	unconstrained	$1 + G[pq - q(q-1)/2] + Gp$
CUUC	constrained	unconstrained	unconstrained	constrained	$1 + G[pq - q(q-1)/2] + G$
CUCU	constrained	unconstrained	constrained	unconstrained	$1 + G[pq - q(q-1)/2] + p$
CUCU	constrained	unconstrained	constrained	constrained	$1 + G[pq - q(q-1)/2] + 1$
CCUU	constrained	constrained	unconstrained	unconstrained	$1 + [pq - q(q-1)/2] + Gp$
CCUC	constrained	constrained	unconstrained	constrained	$1 + [pq - q(q-1)/2] + G$
CCCU	constrained	constrained	constrained	unconstrained	$1 + [pq - q(q-1)/2] + p$
CCCC	constrained	constrained	constrained	constrained	$1 + [pq - q(q-1)/2] + 1$

the parameters in (7); the fitted model classifies each of the $n - m$ unlabeled observations through the corresponding maximum *a posteriori* probability (MAP). As a special case, if $m = 0$, we obtain the clustering scenario. Drawing on Hosmer Jr. (1973), Titterton et al. (1985, Section 4.3.3) point out that knowing the label of just a small proportion of observations *a priori* can lead to improved clustering performance.

Notationally, if the i th observation is labeled, denote with $\tilde{\mathbf{z}}_i = (\tilde{z}_{i1}, \dots, \tilde{z}_{iG})$ its component membership indicator. Then, arranging the data so that the first m observations are labeled, the complete-data likelihood becomes

$$\begin{aligned}
L_c(\boldsymbol{\theta}) &= \prod_{i=1}^m \prod_{g=1}^G \left[\pi_g \phi(y_i | \mathbf{x}_i; \mu(\mathbf{x}; \boldsymbol{\beta}_g), \sigma_g^2) \phi(\mathbf{x}_i | \mathbf{u}_i; \boldsymbol{\mu}_g, \boldsymbol{\Lambda}_g, \boldsymbol{\Psi}_g) \phi(\mathbf{u}_{ig}) \right]^{\tilde{z}_{ig}} \\
&\times \prod_{i=m+1}^n \prod_{g=1}^G \left[\pi_g \phi(y_i | \mathbf{x}_i; \mu(\mathbf{x}; \boldsymbol{\beta}_g), \sigma_g^2) \phi(\mathbf{x}_i | \mathbf{u}_i; \boldsymbol{\mu}_g, \boldsymbol{\Lambda}_g, \boldsymbol{\Psi}_g) \phi(\mathbf{u}_{ig}) \right]^{z_{ig}}.
\end{aligned}$$

For notational convenience, in this paper we prefer to present the AECM algorithm in the model-based clustering paradigm (cf. Section 4). However, the extension to the model-based classification context is simply obtained by substituting the ‘dynamic’ (with respect to the iterations of the algorithm) $\mathbf{z}_1, \dots, \mathbf{z}_m$ with the ‘static’ $\tilde{\mathbf{z}}_1, \dots, \tilde{\mathbf{z}}_m$.

4 Parameter Estimation

4.1 The AECM algorithm

The AECM algorithm (Meng and van Dyk, 1997) is used for fitting all the models within the family defined in Section 1. This algorithm is an extension of the expectation-maximization (EM) algorithm (Dempster et al., 1977) that uses different specifications of missing data at each stage. Let $\mathcal{S} = \{(\mathbf{x}'_i, y_i)'; i = 1, \dots, n\}$ be a sample of size n from (7). In the EM framework, the generic observation $(\mathbf{x}'_i, y_i)'$ is viewed as being incomplete; its complete counterpart is given by $(\mathbf{x}'_i, y_i, \mathbf{u}'_{ig}, \mathbf{z}'_i)'$, where \mathbf{z}_i is the component-label vector in which $z_{ig} = 1$ if $(\mathbf{x}'_i, y_i)'$ comes from Ω_g and $z_{ig} = 0$ otherwise. Then the complete-data likelihood, by considering the result in (5), can be written as

$$L_c(\boldsymbol{\theta}) = \prod_{i=1}^n \prod_{g=1}^G \left[\pi_g \phi(y_i | \mathbf{x}_i; \mu(\mathbf{x}; \boldsymbol{\beta}_g), \sigma_g^2) \phi(\mathbf{x}_i | \mathbf{u}_i; \boldsymbol{\mu}_g, \boldsymbol{\Lambda}_g \boldsymbol{\Lambda}'_g + \boldsymbol{\Psi}_g) \phi(\mathbf{u}_{ig}) \right]^{z_{ig}}.$$

The idea of the AECM algorithm is to partition $\boldsymbol{\theta}$, say $\boldsymbol{\theta}_1 = (\boldsymbol{\theta}'_1, \boldsymbol{\theta}'_2)'$, in such a way that the likelihood is easy to maximize for $\boldsymbol{\theta}_1$ given $\boldsymbol{\theta}_2$ and *vice versa*. The AECM algorithm consists of two cycles, each containing an E-step and a CM-step. The two CM-steps correspond to the partition of $\boldsymbol{\theta}$ into $\boldsymbol{\theta}_1$ and $\boldsymbol{\theta}_2$. Then, we can iterate between these two conditional maximizations until convergence. In the next two sections, we illustrate the two cycles for the UUUU model only. Details on the other models of the family are given in Appendix B.

4.2 First cycle

Here, $\boldsymbol{\theta}_1 = \{\pi_g, \boldsymbol{\beta}_g, \boldsymbol{\mu}_g, \sigma_g^2; g = 1, \dots, G\}$, where the missing data are the unobserved group labels \mathbf{z}_i , $i = 1, \dots, n$. The complete-data likelihood is

$$L_1(\boldsymbol{\theta}_1) = \prod_{i=1}^n \prod_{g=1}^G \left[\pi_g \phi(y_i | \mathbf{x}_i; \mu(\mathbf{x}_i; \boldsymbol{\beta}_g), \sigma_g^2) \phi(\mathbf{x}_i; \boldsymbol{\mu}_g, \boldsymbol{\Sigma}_g) \right]^{z_{ig}}.$$

Consider the complete-data log-likelihood

$$\begin{aligned}
l_{c1}(\boldsymbol{\theta}_1) &= \sum_{i=1}^n \sum_{g=1}^G z_{ig} \ln \left[\pi_g \phi(y_i | \mathbf{x}_i; \mu(\mathbf{x}_i; \boldsymbol{\beta}_g), \sigma_g^2) \phi(\mathbf{x}_i; \boldsymbol{\mu}_g, \boldsymbol{\Lambda}_g, \boldsymbol{\Psi}_g) \right] \\
&= -\frac{n(p+1)}{2} \ln 2\pi - \frac{1}{2} \sum_{i=1}^n \sum_{g=1}^G z_{ig} \ln \sigma_g^2 - \frac{1}{2} \sum_{i=1}^n \sum_{g=1}^G z_{ig} \frac{(y_i - \beta_{0g} - \boldsymbol{\beta}'_{1g} \mathbf{x}_i)^2}{\sigma_g^2} + \\
&\quad - \frac{1}{2} \sum_{i=1}^n \sum_{g=1}^G z_{ig} \ln |\boldsymbol{\Sigma}_g| - \frac{1}{2} \sum_{i=1}^n \sum_{g=1}^G z_{ig} (\mathbf{x}_i - \boldsymbol{\mu}_g)' \boldsymbol{\Sigma}_g^{-1} (\mathbf{x}_i - \boldsymbol{\mu}_g) + \sum_{g=1}^G n_g \ln \pi_g,
\end{aligned}$$

where $n_g = \sum_{i=1}^n z_{ig}$. Because $\boldsymbol{\Sigma}_g = \boldsymbol{\Lambda} \boldsymbol{\Lambda}'_g + \boldsymbol{\Psi}_g$, we get

$$\begin{aligned}
l_{c1}(\boldsymbol{\theta}_1) &= -\frac{n(p+1)}{2} \ln 2\pi - \frac{1}{2} \sum_{i=1}^n \sum_{g=1}^G z_{ig} \ln \sigma_g^2 \\
&\quad - \frac{1}{2} \sum_{i=1}^n \sum_{g=1}^G z_{ig} \frac{(y_i - \beta_{0g} - \boldsymbol{\beta}'_{1g} \mathbf{x}_i)^2}{\sigma_g^2} - \frac{1}{2} \sum_{i=1}^n \sum_{g=1}^G z_{ig} \ln |\boldsymbol{\Lambda} \boldsymbol{\Lambda}'_g + \boldsymbol{\Psi}_g| \\
&\quad - \frac{1}{2} \sum_{i=1}^n \sum_{g=1}^G z_{ig} \text{tr} \left\{ (\mathbf{x}_i - \boldsymbol{\mu}_g) (\mathbf{x}_i - \boldsymbol{\mu}_g)' (\boldsymbol{\Lambda}_g \boldsymbol{\Lambda}'_g + \boldsymbol{\Psi}_g)^{-1} \right\} + \sum_{g=1}^G n_g \ln \pi_g.
\end{aligned}$$

The E-step on the first cycle of the $(k+1)$ st iteration requires the calculation of $Q_1(\boldsymbol{\theta}_1; \boldsymbol{\theta}^{(k)}) = \mathbb{E}_{\boldsymbol{\theta}^{(k)}} [l_c(\boldsymbol{\theta}_1) | \mathcal{S}]$, which is the expected complete-data log-likelihood given the observed data and using the current estimate $\boldsymbol{\theta}^{(k)}$ for $\boldsymbol{\theta}$. In practice, it requires calculating $\mathbb{E}_{\boldsymbol{\theta}^{(k)}} [Z_{ig} | \mathcal{S}]$; this step is achieved by replacing each z_{ig} by $z_{ig}^{(k+1)}$, where

$$z_{ig}^{(k+1)} = \frac{\pi_j^{(k)} \phi(y_i | \mathbf{x}_i; \mu(\mathbf{x}_i; \boldsymbol{\beta}_g^{(k)}), \sigma_g^{2(k)}) \phi(\mathbf{x}_i | \boldsymbol{\mu}_g^{(k)}, \boldsymbol{\Lambda}_g^{(k)}, \boldsymbol{\Psi}_g^{(k)})}{\sum_{j=1}^G \pi_j^{(k)} \phi(y_i | \mathbf{x}_i; \mu(\mathbf{x}_i; \boldsymbol{\beta}_j^{(k)}), \sigma_j^{2(k)}) \phi(\mathbf{x}_i | \boldsymbol{\mu}_j^{(k)}, \boldsymbol{\Lambda}_j^{(k)}, \boldsymbol{\Psi}_j^{(k)})}.$$

For the M-step, the maximization of this complete-data log-likelihood yields

$$\begin{aligned}
\pi_g^{(k+1)} &= \frac{1}{n} \sum_{i=1}^n z_{ig}^{(k+1)} \\
\boldsymbol{\mu}_g^{(k+1)} &= \frac{1}{n_g} \sum_{i=1}^n z_{ig}^{(k+1)} \mathbf{x}_i \\
\boldsymbol{\beta}_{1g}^{(k+1)} &= \left[\frac{1}{n_g} \sum_{i=1}^n z_{ig}^{(k+1)} y_i (\mathbf{x}_i - \boldsymbol{\mu}_g^{(k+1)}) \right] \left[\frac{1}{n_g} \sum_{i=1}^n z_{ig}^{(k+1)} \mathbf{x}_i' \mathbf{x}_i - \boldsymbol{\mu}_g'^{(k+1)} \boldsymbol{\mu}_g^{(k+1)} \right]^{-1} \\
\beta_{0g}^{(k+1)} &= \frac{1}{n_g} \sum_{i=1}^n z_{ig}^{(k+1)} y_i - \boldsymbol{\beta}_{1g}'^{(k+1)} \boldsymbol{\mu}_g^{(k+1)} \\
\sigma_g^{2(k+1)} &= \frac{1}{n_g} \sum_{i=1}^n z_{ig}^{(k+1)} \left\{ y_i - \left(\beta_{0g}^{(k+1)} + \boldsymbol{\beta}_{1g}'^{(k+1)} \mathbf{x}_i \right) \right\}^2,
\end{aligned}$$

where $n_g^{(k+1)} = \sum_{i=1}^n z_{ig}^{(k+1)}$. Following the notation in McLachlan and Peel (2000), we set $\boldsymbol{\theta}^{(k+1/2)} = \{\boldsymbol{\theta}_1^{(k+1)}, \boldsymbol{\theta}_2^{(k)}\}$.

4.3 Second cycle

Here, $\boldsymbol{\theta}_2 = \{\boldsymbol{\Sigma}_g; g = 1, \dots, G\} = \{\boldsymbol{\Lambda}_g, \boldsymbol{\Psi}_g; g = 1, \dots, G\}$, where the missing data are the unobserved group labels \mathbf{z}_i and the latent factors \mathbf{u}_{ig} , $i = 1, \dots, n$ and $g = 1, \dots, G$. Therefore, the complete-data likelihood is

$$\begin{aligned}
L_{c2}(\boldsymbol{\theta}_2) &= \prod_{i=1}^n \prod_{g=1}^G \left[\phi(y_i | \mathbf{x}_i, \mathbf{u}_{ig}; \boldsymbol{\mu}(\mathbf{x}_i; \boldsymbol{\beta}_g^{(k+1)}), \sigma_g^{2(k+1)}) \phi(\mathbf{x}_i | \mathbf{u}_{ig}; \boldsymbol{\mu}_g^{(k+1)}, \boldsymbol{\Sigma}_g) \phi(\mathbf{u}_{ig}) \pi_g^{(k+1)} \right]^{z_{ig}} \\
&= \prod_{i=1}^n \prod_{g=1}^G \left[\phi(y_i | \mathbf{x}_i; \boldsymbol{\mu}(\mathbf{x}_i; \boldsymbol{\beta}_g^{(k+1)}), \sigma_g^{2(k+1)}) \phi(\mathbf{x}_i | \mathbf{u}_{ig}; \boldsymbol{\mu}_g^{(k+1)}, \boldsymbol{\Lambda}_g, \boldsymbol{\Psi}_g) \phi(\mathbf{u}_{ig}) \pi_g^{(k+1)} \right]^{z_{ig}},
\end{aligned}$$

because Y is conditionally independent of \mathbf{U} given $\mathbf{X} = \mathbf{x}$ and

$$\begin{aligned}
\phi(\mathbf{x}_i | \mathbf{u}_{ig}; \boldsymbol{\mu}_g^{(k+1)}, \boldsymbol{\Psi}_g) &= \frac{1}{|2\pi\boldsymbol{\Psi}_g|^{1/2}} \exp \left\{ -\frac{1}{2} (\mathbf{x}_i - \boldsymbol{\mu}_g^{(k+1)} - \boldsymbol{\Lambda}_g \mathbf{u}_{ig})' \boldsymbol{\Psi}_g^{-1} (\mathbf{x}_i - \boldsymbol{\mu}_g^{(k+1)} - \boldsymbol{\Lambda}_g \mathbf{u}_{ig}) \right\} \\
\phi(\mathbf{u}_{ig}) &= \frac{1}{(2\pi)^{q/2}} \exp \left\{ -\frac{1}{2} \mathbf{u}_{ig}' \mathbf{u}_{ig} \right\}.
\end{aligned}$$

Hence, the complete-data log-likelihood is

$$\begin{aligned}
l_{c2}(\boldsymbol{\theta}_2) = & -\frac{n(p+q+1)}{2} \ln(2\pi) - \frac{1}{2} \sum_{i=1}^n \sum_{g=1}^G z_{ig} \ln \sigma_g^{2(k+1)} + \\
& -\frac{1}{2} \sum_{i=1}^n \sum_{g=1}^G z_{ig} \frac{(y_i - \beta_{0g}^{(k+1)} - \boldsymbol{\beta}_{1g}^{'(k+1)} \mathbf{x}_i)^2}{2\hat{\sigma}_g^2} + \sum_{g=1}^G n_g \ln \pi_g + \frac{1}{2} \sum_{i=1}^n \sum_{g=1}^G z_{ig} \ln |\boldsymbol{\Psi}_g^{-1}| + \\
& -\frac{1}{2} \sum_{i=1}^n \sum_{g=1}^G z_{ig} \text{tr} \left\{ \left(\mathbf{x}_i - \boldsymbol{\mu}_g^{(k+1)} - \boldsymbol{\Lambda}_g \mathbf{u}_{ig} \right) \left(\mathbf{x}_i - \boldsymbol{\mu}_g^{(k+1)} - \boldsymbol{\Lambda}_g \mathbf{u}_{ig} \right)' \boldsymbol{\Psi}_g^{-1} \right\},
\end{aligned}$$

where we set

$$\mathbf{S}_g^{(k+1)} = \frac{1}{n_g^{(k+1)}} \sum_{i=1}^n z_{ig}^{(k+1)} \left(\mathbf{x}_i - \boldsymbol{\mu}_g^{(k+1)} \right) \left(\mathbf{x}_i - \boldsymbol{\mu}_g^{(k+1)} \right)'.$$

The E-step on the second cycle of the $(k+1)$ st iteration requires the calculation of $Q_2(\boldsymbol{\theta}_2; \boldsymbol{\theta}^{(k+1/2)}) = \mathbb{E}_{\boldsymbol{\theta}^{(k+1/2)}} [l_{c2}(\boldsymbol{\theta}_2) | \mathcal{S}]$. Therefore, we must calculate the following conditional expectations: $\mathbb{E}_{\boldsymbol{\theta}^{(k+1/2)}} (Z_{ig} | \mathcal{S})$, $\mathbb{E}_{\boldsymbol{\theta}^{(k+1/2)}} (Z_{ig} \mathbf{U}_{ig} | \mathcal{S})$, and $\mathbb{E}_{\boldsymbol{\theta}^{(k+1/2)}} (Z_{ig} \mathbf{U}_{ig} \mathbf{U}_{ig}' | \mathcal{S})$. Based on (2.2), these are given by

$$\begin{aligned}
\mathbb{E}_{\boldsymbol{\theta}^{(k+1/2)}} (Z_{ig} \mathbf{U}_{ig} | \mathcal{S}) &= z_{ig}^{(k+1)} \boldsymbol{\gamma}_g^{(k)} \left(\mathbf{x}_i - \boldsymbol{\mu}_g^{(k+1)} \right) \\
\mathbb{E}_{\boldsymbol{\theta}^{(k+1/2)}} (Z_{ig} \mathbf{U}_{ig} \mathbf{U}_{ig}' | \mathcal{S}) &= z_{ig}^{(k+1)} \left\{ \mathbf{I}_q - \boldsymbol{\gamma}_g^{(k)} \boldsymbol{\Lambda}_g^{(k)} + \boldsymbol{\gamma}_g^{(k)} \mathbf{S}_g \boldsymbol{\gamma}_g^{'(k)} \right\} = z_{ig}^{(k+1)} \boldsymbol{\Theta}_g^{(k)},
\end{aligned}$$

where

$$\boldsymbol{\gamma}_g^{(k)} = \boldsymbol{\Lambda}_g^{'(k)} \left(\boldsymbol{\Lambda}_g^{(k)} \boldsymbol{\Lambda}_g^{'(k)} + \boldsymbol{\Psi}_g^{(k)} \right)^{-1} \quad (8)$$

$$\boldsymbol{\Theta}_g^{(k)} = \mathbf{I}_q - \boldsymbol{\gamma}_g^{(k)} \boldsymbol{\Lambda}_g^{(k)} + \boldsymbol{\gamma}_g^{(k)} \mathbf{S}_g^{(k+1)} \boldsymbol{\gamma}_g^{'(k)}. \quad (9)$$

Thus, the g th term of the expected complete-data log-likelihood $Q_2(\boldsymbol{\theta}_2; \boldsymbol{\theta}^{(k+1/2)})$ becomes

$$\begin{aligned}
Q_2(\boldsymbol{\Lambda}_g, \boldsymbol{\Psi}_g; \boldsymbol{\theta}^{(k+1/2)}) = & C(\boldsymbol{\theta}_1^{(k+1)}) + \frac{1}{2} n_g^{(k+1)} \ln |\boldsymbol{\Psi}_g^{-1}| - \frac{1}{2} n_g^{(k+1)} \text{tr} \left\{ \mathbf{S}_g^{(k+1)} \boldsymbol{\Psi}_g^{-1} \right\} \\
& + n_g^{(k+1)} \text{tr} \left\{ \boldsymbol{\Lambda}_g \boldsymbol{\gamma}_g^{(k)} \mathbf{S}_g^{(k+1)} \boldsymbol{\Psi}_g^{-1} \right\} - \frac{1}{2} n_g^{(k+1)} \text{tr} \left\{ \boldsymbol{\Lambda}_g' \boldsymbol{\Psi}_g^{-1} \boldsymbol{\Lambda}_g \boldsymbol{\Theta}_g^{(k)} \right\},
\end{aligned} \quad (10)$$

where $C(\boldsymbol{\theta}_1^{(k+1)})$ denotes the terms in (4.3) that do not depend on $\boldsymbol{\theta}_2$. Then

(10) is maximized for $\{\hat{\Lambda}, \hat{\Psi}\}$, satisfying

$$\begin{aligned}\frac{\partial Q_2}{\partial \Lambda_g} &= n_g^{(k+1)} \Psi_g^{-1} S_g^{(k+1)} \gamma_g'^{(k)} - n_g^{(k+1)} \Psi_g^{-1} \Lambda_g \Theta_g^{(k)} = \mathbf{0} \\ \frac{\partial Q_2}{\partial \Psi_g^{-1}} &= \frac{1}{2} n_g^{(k+1)} \Psi_g - \frac{1}{2} n_g^{(k+1)} S_g^{(k+1)} + n_g^{(k+1)} S_g'^{(k+1)} \gamma_g'^{(k)} \Lambda_g' - \frac{1}{2} n_g^{(k+1)} \Lambda_g \Theta_g^{(k)} \Lambda_g' = \mathbf{0}.\end{aligned}$$

Therefore,

$$S_g^{(k+1)} \gamma_g'^{(k)} - \Lambda_g \Theta_g^{(k)} = \mathbf{0} \quad (11)$$

$$\Psi_g - S_g^{(k+1)} + 2S_g'^{(k+1)} \gamma_g'^{(k)} \Lambda_g' - \Lambda_g \Theta_g^{(k)} \Lambda_g' = \mathbf{0}. \quad (12)$$

From (11), we get

$$\hat{\Lambda}_g = S_g^{(k+1)} \gamma_g'^{(k)} \Theta_g^{-1}, \quad (13)$$

and substituting in (12) we get

$$\Psi_g - S_g^{(k+1)} + 2S_g^{(k+1)} \gamma_g'^{(k)} \left(S_g^{(k+1)} \gamma_g'^{(k)} \Theta_g^{-1} \right)' - \left(S_g \hat{\gamma}_g' \Theta_g^{-1} \right) \Theta_g \left(S_g \hat{\gamma}_g' \Theta_g^{-1} \right)' = \mathbf{0}$$

which yields

$$\hat{\Psi}_g = \text{diag} \left\{ S_g^{(k+1)} - \hat{\Lambda}_g \hat{\gamma}_g S_g^{(k+1)} \right\}. \quad (14)$$

Hence, the maximum likelihood estimates for Λ and Ψ are obtained by iteratively computing

$$\begin{aligned}\Lambda_g^+ &= S_g^{(k+1)} \gamma_g' \Theta_g^{-1} \\ \Psi_g^+ &= \text{diag} \left\{ S_g^{(k+1)} - \Lambda_g^+ \gamma_g S_g^{(k+1)} \right\},\end{aligned}$$

where the superscript $+$ denotes the update estimate. Using (9) and (9), we get

$$\begin{aligned}\gamma_g^+ &= \Lambda_g'^+ \left(\Lambda_g^+ \Lambda_g'^+ + \Psi_g^+ \right)^{-1} \\ \Theta_g^+ &= \mathbf{I}_q - \gamma_g^+ \Lambda_g^+ + \gamma_g^+ S_g^{(k+1)} \gamma_g'^+.\end{aligned} \quad (15)$$

4.4 Outline of the algorithm

In summary, the procedure can be described as follows. For a given initial guess $\theta^{(0)}$, on the $(k+1)$ st iteration, the algorithm carries out the following steps for $g = 1, \dots, G$:

- (1) Compute $\pi_g^{(k+1)}, \mu_g^{(k+1)}, \beta_g^{(k+1)}, \sigma_g^{2(k+1)}$;
- (2) Set $\Lambda_g \leftarrow \Lambda_g^{(k)}$ and $\Psi \leftarrow \Psi_g^{(k)}$, and compute γ_g and Θ_g ;

- (3) Repeat the following steps until convergence on Λ_g and Ψ_g :
- (a) Set $\Lambda_g^+ \leftarrow \mathbf{S}_g^{(k+1)} \gamma_g' \Theta_g^{-1}$ and $\Psi_g^+ \leftarrow \text{diag} \left\{ \mathbf{S}_g^{(k+1)} - \Lambda_g^+ \gamma_g \mathbf{S}_g^{(k+1)} \right\}$;
 - (b) Set $\gamma_g^+ \leftarrow \Lambda_g'^+ \left(\Lambda_g^+ \Lambda_g'^+ + \Psi_g^+ \right)^{-1}$ and $\Theta_g^+ \leftarrow \mathbf{I}_q - \gamma_g^+ \Lambda_g^+ + \gamma_g^+ \mathbf{S}_g^{(k+1)} \gamma_g'^+$;
 - (c) Set $\Lambda_g \leftarrow \Lambda_g^+$, $\Psi_g \leftarrow \Psi_g^+$, $\gamma_g \leftarrow \gamma_g^+$, and $\Theta_g \leftarrow \Theta_g^+$.

4.5 AECM initialization: a 5-step procedure

The choice of starting values is a well known and important issue with respect to EM-based algorithms. The standard approach consists of selecting a value for $\theta^{(0)}$. An alternative method, more natural in the authors' opinion, consists of choosing a value for $\mathbf{z}_i^{(0)}$, $i = 1, \dots, n$ (see McLachlan and Peel, 2000, p. 54). Within this approach, and due to the hierarchical structure of the CWFA family of parsimonious models, we propose a 5-step hierarchical initialization procedure.

For a fixed number of groups G , let $\mathbf{z}_i^{(0)}$, $i = 1, \dots, n$, be the initial classification for the AECM algorithm, so that $z_{ig}^{(0)} \in \{0, 1\}$ and $\sum_g z_{ig}^{(0)} = 1$. The set $\{\mathbf{z}_i^{(0)}; i = 1, \dots, n\}$ can be obtained either through some clustering procedure (here we consider the k -means method) or by random initialization, for example by sampling from a multinomial distribution with probabilities $(1/G, \dots, 1/G)$. Then, at the first step of the procedure, the most constrained CCCC model is estimated from these starting values. At the second step, the resulting (AECM-estimated) \hat{z}_{ig} are taken as the starting group membership labels to initialize the AECM-algorithm of the four models $\{\text{UCCC}, \text{CUCC}, \text{CCUC}, \text{CCCU}\}$ obtained by relaxing one of the four constraints. At the third step, the AECM-algorithm for each of the six models $\{\text{CCUU}, \text{CUCU}, \text{UCCU}, \text{CUUC}, \text{UCUC}, \text{UCCC}\}$ with two constraints is initialized using the \hat{z}_{ig} from the previous step and the model with the highest likelihood. For example, to initialize CCUU we use the \hat{z}_{ig} from the model having the highest likelihood between CCCU and CCUC. In this fashion, the initialization procedure continues according to the scheme displayed in Fig. 1, until the least constrained model UUUU is estimated at the fifth step.

For all of the models in the CWFA family, in analogy with McNicholas and Murphy (2008), the initial values for the elements of Λ_g and Ψ_g are generated from the eigen-decomposition of \mathbf{S}_g as follows. The \mathbf{S}_g are computed based on the values of $\mathbf{z}_{ig}^{(0)}$. The eigen-decomposition of each \mathbf{S}_g is obtained using the Householder reduction and the QL method (details given by Press et al., 1992). Then the initial values of the elements of Λ_g are set as $\lambda_{ij} = \sqrt{d_j} \rho_{ij}$, where d_j is the j th largest eigenvalue of \mathbf{S}_g and ρ_{ij} is the i th element of the eigenvector corresponding to the j th largest eigenvalue of \mathbf{S}_g , where $i \in \{1, 2, \dots, d\}$ and $j \in \{1, 2, \dots, q\}$. The Ψ_g are then initialized as $\Psi_g = \text{diag} \left(\mathbf{S}_g - \Lambda_g \Lambda_g' \right)$.




Fig. 1. Relationships among the models in the 5-step hierarchical initialization procedure. Arrows are oriented from the model used to initialize to the model to be estimated.

4.6 Convergence criterion

The Aitken acceleration procedure (Aitken, 1926) is used to estimate the asymptotic maximum of the log-likelihood at each iteration of the AECM algorithm. Based on this estimate, a decision is made about whether the algorithm has reached convergence, i.e., whether the log-likelihood is sufficiently close to its estimated asymptotic value. The Aitken acceleration at iteration k is given by

$$a^{(k)} = \frac{l^{(k+1)} - l^{(k)}}{l^{(k)} - l^{(k-1)}},$$

where $l^{(k+1)}$, $l^{(k)}$, and $l^{(k-1)}$ are the (observed-data) log-likelihood values from iterations $k+1$, k , and $k-1$, respectively. Then, the asymptotic estimate of the log-likelihood at iteration $k+1$ is

$$l_{\infty}^{(k+1)} = l^{(k)} + \frac{1}{1 - a^{(k)}} (l^{(k+1)} - l^{(k)})$$

(Böhning et al., 1994). In the analyses in Section 6, we stop our algorithms when $l_{\infty}^{(k+1)} - l^{(k)} < \epsilon$ (Böhning et al., 1994; McNicholas et al., 2010). Note that we use $\epsilon = 0.05$ for the analyses herein.

5 Model selection and performance assessment

5.1 Model selection

The CWFA model, in addition to $\boldsymbol{\theta}$, is also characterized by the number of latent factors q and by the number of mixture components g . So far, these quantities have been treated as *a priori* fixed. Nevertheless, the estimation of these is required, for practical purposes, when choosing a relevant model.

For model-based clustering and classification, several model selection criteria are used, such as the Bayesian information criterion (BIC; Schwarz, 1978), the integrated completed likelihood (ICL; Biernacki et al., 2000), and the Akaike information criterion (AIC; Sakamoto et al., 1983). Among these, the BIC is the most predominant in the literature and is given by

$$\text{BIC} = 2l(\hat{\boldsymbol{\theta}}) - \eta \ln(n),$$

where $l(\hat{\boldsymbol{\theta}})$ is the (maximized) observed-data log-likelihood and η is the number of free parameters. This is the model selection criterion used in the analyses of Section 6.

5.2 Adjusted Rand index

Although the data analyses of Section 6 are mainly conducted as clustering examples, the true classifications are actually known for these data. In these examples, the adjusted Rand index (ARI; Hubert and Arabie, 1985) is used to measure class agreement. The original Rand Index (RI; Rand, 1971) is based on pairwise comparisons and is obtained by dividing the number of pair agreements (observations that should be in the same group and are, plus those that should not be in the same group and are not) by the total number of pairs. The ARI corrects the RI to account for agreement by chance: a value of ‘1’ indicates perfect agreement, ‘0’ indicates random classification, and negative values indicate a classification that is worse than would be expected by guessing.

6 Data analyses

This section presents the application of the family of parsimonious linear Gaussian models to both artificial and real data sets. Code for the AECM algo-

rithm, described in this paper, was written in the R computing environment (R Development Core Team, 2012).

6.1 Simulated data

6.1.1 Example 1

The first data set consists of a sample of size $n = 175$ drawn from model UUCU with $G = 2$, $n_1 = 75$, $n_2 = 100$, $d = 5$, and $q = 2$ (see Fig. 2 for details). The parameters used for the simulation of the data are given in Table 2 (see

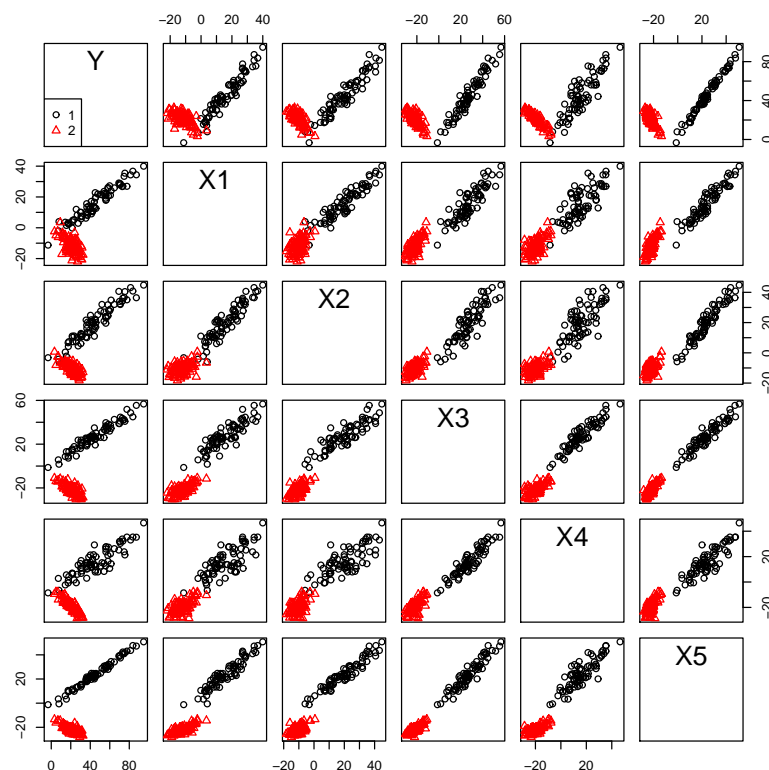


Fig. 2. Example 6.1.1: scatterplot matrix of the simulated data.

Appendix C.1 for details on the covariance matrices Σ_g , $g = 1, \dots, G$).

All of the sixteen CWFA models were fitted to the data for $G \in \{2, 3\}$ and $q \in \{1, 2\}$, resulting in a total of 64 models. As noted above (Section 4.5), initialization of the z_i , $i = 1, \dots, n$, for the most constrained model (CCCC), and for each combination (G, q) , was done using the k -means algorithm according to the `kmeans` function of the R package `stats`. The remaining 15 models, for each combination (G, q) , were initialized using the 5-step hierarchical initialization procedure described in Section 4.5. The BIC values for

Table 2

True and estimated parameters for the simulated data of Example 1.

(a) Means of \mathbf{X}										
Group	$\boldsymbol{\mu}_g$					$\hat{\boldsymbol{\mu}}_g$				
	X_1	X_2	X_3	X_4	X_5	X_1	X_2	X_3	X_4	X_5
1	14.00	18.00	25.00	14.00	22.00	15.88	19.94	27.48	15.81	23.93
2	-12.00	-10.00	-22.00	-20.00	-22.00	-11.95	-10.36	-22.00	-19.67	-22.03

(b) Slopes										
Group	$\boldsymbol{\beta}_{1g}$					$\hat{\boldsymbol{\beta}}_{1g}$				
	X_1	X_2	X_3	X_4	X_5	X_1	X_2	X_3	X_4	X_5
1	0.47	0.02	0.42	0.03	0.87	0.50	0.03	0.46	0.02	0.81
2	-0.02	-0.63	-0.05	-0.85	-0.03	-0.04	-0.57	-0.01	-0.85	-0.18

(c) Conditional std. deviations			(d) Intercepts		
Group	σ_g	$\hat{\sigma}_g$	Group	β_{0g}	$\hat{\beta}_{0g}$
1	2.00	1.24	1	4.50	4.34
2	4.00	3.79	2	-4.20	-6.35

all 64 models were computed and the model with the largest BIC value was selected as the best model. In this example, the model corresponding to the largest BIC value (-5845.997) was a two component ($G = 2$) UUCU model with two latent factors ($q = 2$), the same as the model used to generate the data. The selected model gave a perfect classification and the estimated parameters were very close to the parameters used for data simulation (see Table 2 and Appendix C.1).

Fig. 3 shows the BIC values of all 64 models sorted in an increasing order, where numbers denote the selected number G of groups and colours denote the number q of latent factors. The horizontal line separates the models with a BIC value within 1% of the maximum (over the 64 models) BIC value (hereafter simply referred to as the ‘1% line’). This graphical representation will be referred to as the ‘group-factor plot’ of BIC values. Here, as mentioned earlier, the model with the largest BIC was UUCU (with $G = 2$ and $q = 2$). The subsequent two models, those above the 1% line, were UUUU with $G = 2$ and $q = 2$ (BIC equal to -5867.006) and CUCU with $G = 2$ and $q = 2$ (BIC equal to -5869.839). These two models are structurally very close to the true UUCU model and also yielded perfect classification. It should also be noted




Fig. 3. Group-factor plot of BIC values sorted in an increasing order for Example 1. Numbers denote the selected number G of groups and colours denote the number q of latent factors (red: $q = 1$, blue: $q = 2$). The green number indicates the true model.

that most of the models with high BIC values have $G = 2$ and $q = 2$.

6.1.2 Example 2

For the second data set, a sample of size $n = 235$ was drawn from the CUUC model with $G = 3$ groups (of size $n_1 = 75$, $n_2 = 100$, and $n_3 = 60$) and $q = 2$ latent factors (see Fig. 4).

All 16 CWFA models were fitted to the data for $G \in \{2, 3, 4\}$ and $q \in \{1, 2\}$, resulting in 96 different models. The algorithm was initialized in the same way as for Example 2. The model with the highest BIC (-6579.116) was CUUC with $G = 3$ and $q = 2$, resulting in a perfect classification. The estimated parameters of this model were very close to the true ones (Table 3 and Appendix C.2).

Fig. 5 shows the group-factor plot of BIC values for all 96 models. The other three models above the 1% line are UUUC (BIC= -6583.692), CUUU (BIC= -6637.222), and UUUU (BIC= -6641.798), all with $G = 3$ and $q = 2$. Thus, these models are congruent, with respect to the true one, in terms of G and q . Moreover, they had a covariance structure more similar to the true one (CUUC) and also yielded perfect classification.

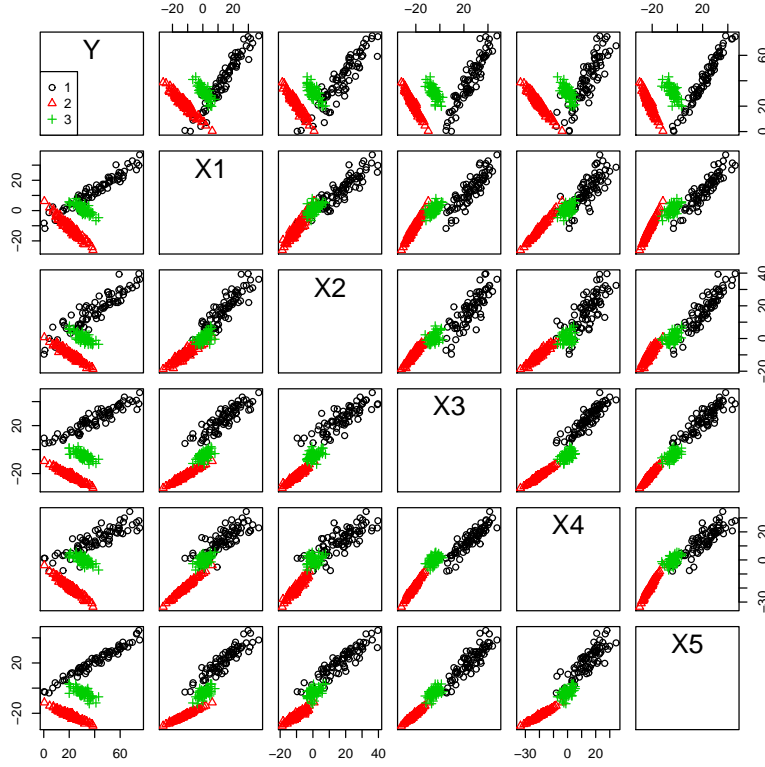


Fig. 4. Scatterplot matrix of the simulated data for Example 2.

6.2 The *f.voles* data set

In addition to the simulated data analyses of Section 6.1, the family of CW-FAs was also applied to a real data set for both clustering and classification purposes.

The *f.voles* data set, detailed in Flury (1997, Table 5.3.7) and available in the *Flury* package for R, consists of measurements of female voles from two species, *M. californicus* and *M. ochrogaster*. The data consist of 86 observations for which we have a binary variable **Species** denoting the species ($n_1 = 45$ *Microtus ochrogaster* and $n_1 = 41$ *M. californicus*), a variable **Age** measured in days, and six remaining variables related to skull measurements. The names of the variables are the same as in the original analysis of this data set by Airoldi and Hoffmann (1984): L_2 = condylo-incisive length, L_9 = length of incisive foramen, L_7 = alveolar length of upper molar tooth row, B_3 = zygomatic width, B_4 = interorbital width, and H_1 = skull height. All of the variables related to the skull are measured in units of 0.1 mm.

The purpose of Airoldi and Hoffmann (1984) was to study age variation in *M. californicus* and *M. ochrogaster* and to predict age on the basis of the

Table 3

True and estimated parameters for the simulated data of Example 2.

(a) Means of \mathbf{X}

Group	$\boldsymbol{\mu}_g$					$\hat{\boldsymbol{\mu}}_g$				
	X_1	X_2	X_3	X_4	X_5	X_1	X_2	X_3	X_4	X_5
1	0.00	0.00	-5.00	0.00	-4.00	0.82	0.48	-5.09	-0.21	-3.75
2	14.00	18.00	25.00	14.00	22.00	13.64	17.44	25.44	14.25	21.44
3	-12.00	-10.00	-22.00	-20.00	-22.00	-12.33	-10.22	-22.25	-20.24	-22.21

(b) Slopes

Group	$\boldsymbol{\beta}_{1g}$					$\hat{\boldsymbol{\beta}}_{1g}$				
	X_1	X_2	X_3	X_4	X_5	X_1	X_2	X_3	X_4	X_5
1	-0.41	-0.87	-0.22	-0.62	-0.06	-0.34	-0.82	-0.32	-0.66	-0.09
2	0.47	0.02	0.42	0.03	0.87	0.51	0.00	0.38	0.05	0.84
3	-0.02	-0.63	-0.05	-0.85	-0.03	-0.04	-0.68	-0.36	-0.44	-0.18

(c) Conditional std. deviations

Group	σ_g	$\hat{\sigma}_g$
1	2.00	2.30
2	2.00	2.30
3	2.00	2.30

(d) Intercepts

Group	β_{0g}	$\hat{\beta}_{0g}$
1	30.00	29.39
2	4.50	5.31
3	-4.20	-6.69

skull measurements. For our purpose, we assume the data are unlabelled with respect to **Species** and that our interest is in evaluating clustering and classification using the family of CWFA models as well as comparing the algorithm with some well-established mixture model-based techniques. Therefore, **Age** can be considered the natural Y variable and the $d = 6$ skull measurements can be considered as the \mathbf{X} variable for the CWFA framework.

6.2.1 Clustering

All sixteen linear Gaussian CWFA models were fitted — assuming no known group membership — for $G \in \{2, \dots, 5\}$ components and $q \in \{1, 2, 3\}$ latent factors, resulting in total of 192 different models. The model with the largest BIC value was CCCU with $G = 3$ and $q = 1$, with a BIC of -3837.698 and an ARI of 0.72. Table 4(a) displays the clustering results from this model.




Fig. 5. Group-factor plot of BIC values sorted in increasing order for Example 2. Numbers denote the selected number G of groups and colours denote the number q of latent factors (red: $q = 1$, blue: $q = 2$). The green number indicates the true model.

Furthermore, Table 4(b) and Table 4(c) show, respectively, the clustering re-

Table 4

Clustering of `f.voies` data using three different clustering approaches.

(a) CWFA					(b) PGMM				
TRUE \	Est.	1	2	3	TRUE \	Est.	1	2	3
	TRUE					TRUE			
<i>ochrogaster</i>	24	21	–		<i>ochrogaster</i>	34	9	2	
<i>californicus</i>	–	–	41		<i>californicus</i>	–	–	41	

(c) MCLUST			
TRUE \	Est.	1	2
	TRUE		
<i>ochrogaster</i>	43	2	
<i>californicus</i>	–	41	

sults of the following model-based clustering approaches applied to the vector $(\mathbf{X}', Y)'$:

PGMM: parsimonious latent Gaussian mixture models as described in McNicholas and Murphy (2008), McNicholas (2010), and McNicholas et al. (2010), and estimated via the `pgmmEM` function of the R-package `pgmm` (McNicholas et al., 2011); and **MCLUST:** parsimonious mixtures of Gaussian distributions as described in Banfield and Raftery (1993), Celeux and Govaert (1995), and Fraley and Raftery (2002), and estimated via the `Mclust` function of the R-package `mclust` (see Fraley et al., 2012, for details).

As seen from Table 4, *M. californicus* was classified correctly using all three approaches. Also, *M. ochrogaster* was classified into two sub-clusters using CWFA and PGMM while MCLUST classified it into one cluster. However, the CWFA approach had no misclassifications between the two species but both PGMM and MCLUST misclassified two *M. ochrogaster* as *M. californicus*.




Fig. 6. Group-factor plot of BIC values sorted in an increasing order for the `f.voies` data. Numbers denote the selected number G of groups and colors denote the number q of latent factors (red: $q = 1$, blue: $q = 2$, green: $q = 3$).

Now, we evaluate the group-factor plot of BIC values, for all 192 models, displayed in Fig. 6. Ten models had a BIC above the 1% line; among them, six were characterized by $G = 3$ components and the remaining four by $G = 2$. However, from Fig. 6, the top four models all had three components, which shows that a three component model was not randomly chosen. Airoldi and Hoffmann (1984) mention that some unexplained geographic variation may exist among the voles. However, no covariate was available with such information. Hence, we opted for the scatter plot matrix to evaluate the presence of sub-clusters, see Fig. 7. Here, the scatter plot of the variables B_3

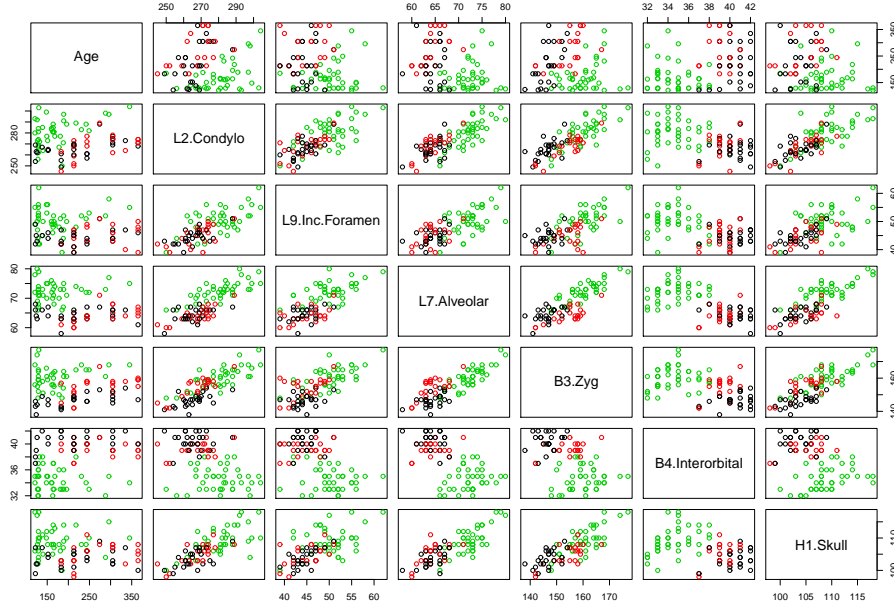


Fig. 7. Scatterplot matrix of `f.voies` data showing the classification observed from CWFA modelling using the clustering framework, where black and red indicates sub-clusters of *M. ochrogaster* species and green indicates *M. californicus* species.

versus B_4 shows the presence of distinct sub-clusters for *M. ochrogaster*, which supports our results attained using CWFA modelling.

6.2.2 Classification

A subset of observations, consisting of 50% of the data, was randomly selected and these observations were assumed to have a known group membership. To allow for the unobserved sub-cluster noted in the clustering application of Section 6.2.1, we ran the algorithm for $G = 2, 3$ and $q = 1, 2, 3$. The best model (CCUU with $G = 2$ and $q = 1$) selected by the BIC (-3843.482) gave a perfect classification, as we can see from Table 5(a).

Table 5

Classification of `f.voies` data assuming that 50% of the observations have known group membership.

(a) 2 known groups				(b) 3 known groups			
TRUE \ Est.	Est.		TRUE	TRUE \ Est.	Est.		
	1	2			1	2	3
<i>ochrogaster</i>	45	–		<i>ochrogaster</i>	28	17	–
<i>californicus</i>	–	41		<i>californicus</i>	–	–	41

We also ran the classification assuming that the data are actually comprised of three known groups. Therefore, using the classification observed by clustering, we also ran the classification algorithm with 50% known (i.e., labelled) and 50% unknown (i.e., unlabelled). To further allow for the unobserved sub-cluster, we ran the algorithm for $G \in \{3, 4\}$ and $q \in \{1, 2, 3\}$. The model selected using the BIC was CCCU with $G = 3$ and $q = 1$, with a BIC value of -3837.383 . Even though the BIC value observed using the classification approach (with three known groups membership) was very close to the BIC value using clustering, the sub-clusters do not have precisely the same classification using the classification and clustering approaches. This could be a consequence of the classification of borderline observations among the sub-clusters using maximum *a posteriori* probability. However, the BIC value for the classification using three known groups was higher than the BIC value using two known groups, which again suggests the presence of sub-clusters.

7 Conclusions, discussion, and future work

In this paper, we introduced a novel family of 16 parsimonious mixture models for model-based clustering and classification. They are linear Gaussian cluster-weighted models in which a latent factor structure is assumed for the explanatory random vector in each mixture component. The parsimonious versions are obtained by combining all of the constraints described in McNicholas and Murphy (2008) with one of the constraints illustrated in Ingrassia, Minotti, and Punzo (2012). Due to the introduction of a latent factor structure, the parameters are linear in dimensionality as opposed to the traditional linear Gaussian CWM where the parameters grow quadratically; therefore, our approach is more suitable for modelling complex high dimensional data. The AECM algorithm (Meng and van Dyk, 1997) was used for maximum likelihood estimation of the model parameters. Being based on the EM algorithm, it is very sensitive to the starting values due to presence of multiple local maxima in high dimensional space. To overcome this problem, we proposed a 5-step hierarchical initialization procedure that utilizes the nested structures of the models within the family. Because these models have a hierarchical/nested structure, this initialization procedure guarantees a natural ranking on the likelihoods of the models in our family. In other words, the procedure restricts a model A, which is nested in a model B, from having a greater likelihood than model B. Using artificial and real data, we demonstrated that these models give very good clustering performance and that the AECM algorithms used were able to recover the parameters very well.

With regard to the latent factor structure, for a latent dimension $q > 1$, the loading matrix $\mathbf{\Lambda}$ is unidentifiable because the model is still satisfied even when the latent factor \mathbf{u}_i is replaced by $\mathbf{H}\mathbf{u}_i$ and $\mathbf{\Lambda}$ by $\mathbf{\Lambda}\mathbf{H}'$, where \mathbf{H} is any orthog-

onal matrix of order q (McLachlan and Peel, 2000). This results in an infinite number of possibilities for $\mathbf{\Lambda}$. Even though this does not affect the clustering algorithm interpretation of the estimated $\mathbf{\Lambda}$ is not informative because $\mathbf{\Lambda}\mathbf{\Lambda}'$ is unique. A future avenue of research is to explore further constraints on the factor loading matrix to ensure a uniquely defined factor loading matrix $\mathbf{\Lambda}$.

Also, while the BIC was able to identify the correct model, the choice of a convenient model selection criterion for these models is still an open question. Some future work will be devoted to the search for good model selection criteria for these models. Finally, here we assumed that the number of factors was the same across groups, which might be too restrictive. However, assuming otherwise also increases the number of models that need to be fitted, resulting in huge computational burden. Approaches such as variational Bayes approximations might be useful for significantly reducing the number of models that need to be fitted.

A The conditional distribution of $Y|\mathbf{x}, \mathbf{u}$

To compute the distribution of $Y|\mathbf{x}, \mathbf{u}$, we begin by recalling that if $\mathbf{Z} \sim N_q(\mathbf{m}, \mathbf{\Gamma})$ is a random vector with values in \mathbb{R}^q and if \mathbf{Z} is partitioned as $\mathbf{Z} = (\mathbf{Z}'_1, \mathbf{Z}'_2)'$, where \mathbf{Z}_1 takes values in \mathbb{R}^{q_1} and \mathbf{Z}_2 in $\mathbb{R}^{q_2} = \mathbb{R}^{q-q_1}$, then we can write

$$\mathbf{m} = \begin{bmatrix} \mathbf{m}_1 \\ \mathbf{m}_2 \end{bmatrix} \quad \text{and} \quad \mathbf{\Gamma} = \begin{bmatrix} \mathbf{\Gamma}_{11} & \mathbf{\Gamma}_{12} \\ \mathbf{\Gamma}_{21} & \mathbf{\Gamma}_{22} \end{bmatrix}.$$

Now, because \mathbf{Z} has a multivariate normal distribution, $\mathbf{Z}_1|\mathbf{Z}_2 = \mathbf{z}_2$ and \mathbf{Z}_2 are statistically independent with $\mathbf{Z}_1|\mathbf{Z}_2 = \mathbf{z}_2 \sim N_{q_1}(\mathbf{m}_{1|2}, \mathbf{\Gamma}_{1|2})$ and $\mathbf{Z}_2 \sim N_{q_2}(\mathbf{m}_2, \mathbf{\Gamma}_{22})$, where

$$\mathbf{m}_{1|2} = \mathbf{m}_1 + \mathbf{\Gamma}_{12}\mathbf{\Gamma}_{22}^{-1}(\mathbf{z}_2 - \mathbf{m}_2) \quad \text{and} \quad \mathbf{\Gamma}_{1|2} = \mathbf{\Gamma}_{11} - \mathbf{\Gamma}_{12}\mathbf{\Gamma}_{22}^{-1}\mathbf{\Gamma}_{21}. \quad (\text{A.1})$$

Therefore, setting $\mathbf{Z} = (\mathbf{Z}'_1, \mathbf{Z}'_2)'$, where $\mathbf{Z}'_1 = Y$ and $\mathbf{Z}'_2 = (\mathbf{X}', \mathbf{U})'$, gives $\mathbf{m}_1 = \beta_0 + \beta'_1\boldsymbol{\mu}$ and $\mathbf{m}_2 = (\boldsymbol{\mu}', \mathbf{0}')'$, with the elements in $\mathbf{\Gamma}$ given by

$$\mathbf{\Gamma}_{11} = \beta'_1\Sigma\beta_1 + \sigma^2, \quad \mathbf{\Gamma}_{22} = \begin{bmatrix} \Sigma & \mathbf{\Lambda} \\ \mathbf{\Lambda}' & \mathbf{I}_q \end{bmatrix}, \quad \text{and} \quad \mathbf{\Gamma}_{12} = \begin{bmatrix} \beta'_1\Sigma & \beta'_1\mathbf{\Lambda} \end{bmatrix}.$$

It follows that $Y|\mathbf{x}, \mathbf{u}$ is Gaussian with mean $\mathbf{m}_{y|\mathbf{x}, \mathbf{u}} = \mathbb{E}(Y|\mathbf{x}, \mathbf{u})$ and variance $\sigma_{y|\mathbf{x}, \mathbf{u}}^2 = \text{Var}(Y|\mathbf{x}, \mathbf{u})$, in accordance with the formulae in (A.1). Because the inverse matrix of $\mathbf{\Gamma}_{22}$ is required in (A.1), the following formula for the inverse

of a partitioned matrix is utilized:

$$\begin{bmatrix} \mathbf{A} & \mathbf{B} \\ \mathbf{C} & \mathbf{D} \end{bmatrix}^{-1} = \begin{bmatrix} (\mathbf{A} - \mathbf{B}\mathbf{D}^{-1}\mathbf{C})^{-1} & -\mathbf{A}^{-1}\mathbf{B}(\mathbf{D} - \mathbf{C}\mathbf{A}^{-1}\mathbf{B})^{-1} \\ -\mathbf{D}^{-1}\mathbf{C}(\mathbf{A} - \mathbf{B}\mathbf{D}^{-1}\mathbf{C})^{-1} & (\mathbf{D} - \mathbf{C}\mathbf{A}^{-1}\mathbf{B})^{-1} \end{bmatrix}.$$

Again, writing $\Sigma = \Lambda\Lambda' + \Psi$, we have

$$\Gamma_{22}^{-1} = \begin{bmatrix} \Sigma & \Lambda \\ \Lambda' & \mathbf{I}_q \end{bmatrix}^{-1} = \begin{bmatrix} \Psi^{-1} & -\Sigma^{-1}\Lambda(\mathbf{I}_q - \Lambda'\Sigma^{-1}\Lambda)^{-1} \\ -\Lambda'\Psi^{-1} & (\mathbf{I}_q - \Lambda'\Sigma^{-1}\Lambda)^{-1} \end{bmatrix}.$$

Moreover, according to the Woodbury identity (Woodbury, 1950):

$$\Sigma^{-1} = (\Lambda\Lambda' + \Psi)^{-1} = \Psi^{-1} - \Psi^{-1}\Lambda(\mathbf{I}_q + \Lambda'\Psi^{-1}\Lambda)^{-1}\Lambda'\Psi^{-1}.$$

Now,

$$\Gamma_{12}\Gamma_{22}^{-1} = \begin{bmatrix} \beta_1'\Sigma & \beta_1'\Lambda \end{bmatrix} \begin{bmatrix} \Psi^{-1} & -\Sigma^{-1}\Lambda(\mathbf{I}_q - \Lambda'\Sigma^{-1}\Lambda)^{-1} \\ -\Lambda'\Psi^{-1} & (\mathbf{I}_q - \Lambda'\Sigma^{-1}\Lambda)^{-1} \end{bmatrix} = \begin{bmatrix} \beta_1' & 0 \end{bmatrix}.$$

Finally, according to (A.1), we have

$$\begin{aligned} \mathbf{m}_{y|x,u} &= \mathbf{m}_1 + \Gamma_{12}\Gamma_{22}^{-1} \begin{bmatrix} \mathbf{z}_2 - \mathbf{m}_2 \end{bmatrix} = (\beta_0 + \beta_1'\boldsymbol{\mu}) + \begin{bmatrix} \beta_1' & 0 \end{bmatrix} \begin{bmatrix} \mathbf{x} - \boldsymbol{\mu} \\ \mathbf{u} - \mathbf{0} \end{bmatrix} = \beta_0 + \beta_1'\mathbf{x}, \\ \sigma_{y|x,u}^2 &= \Gamma_{11} - \Gamma_{12}\Gamma_{22}^{-1}\Gamma_{21} = \beta_1'\Sigma\beta_1 + \sigma^2 - \begin{bmatrix} \beta_1' & 0 \end{bmatrix} \begin{bmatrix} \Sigma\beta_1 \\ \Lambda\beta_1 \end{bmatrix} = \sigma^2. \end{aligned}$$

B Details on the AECM algorithm for the parsimonious models

This appendix details the AECM algorithm for of all the models summarized in Table 1.

B.1 Constraint on the Y variable

In all of the models whose identifier starts with ‘C’, that is the models in which the error variance terms σ_g^2 (of the response variable Y) are constrained

to be equal across groups, i.e., $\sigma_g^2 = \sigma^2$ for $g = 1, \dots, G$, the common variance σ^2 at the $(k+1)$ th iteration of the algorithm is computed as

$$\sigma^{2(k+1)} = \frac{1}{n} \sum_{i=1}^n \sum_{g=1}^G z_{ig}^{(k+1)} \left\{ y_i - \left(\beta_{0g}^{(k+1)} + \beta_{1g}^{\prime(k+1)} \mathbf{x}_i \right) \right\}^2.$$

B.2 Constraints on the \mathbf{X} variable

With respect to the \mathbf{X} variable, as explained in Section 3.2, we considered the following constraints on $\Sigma_g = \Lambda_g \Lambda_g' + \Psi_g$: *i*) equal loading matrices $\Lambda_g = \Lambda$, *ii*) equal error variance $\Psi_g = \Psi$, and *iii*) isotropic assumption: $\Psi_g = \psi_g \mathbf{I}_p$. In such cases, the g th term of the expected complete-data log-likelihood $Q_2(\boldsymbol{\theta}_2; \boldsymbol{\theta}^{(k+1/2)})$, and then the estimates (13) and (14) in Section 4.3, are computed as follows.

B.2.1 Isotropic assumption: $\Psi_g = \psi_g \mathbf{I}_p$

In this case, Equation (10) becomes

$$\begin{aligned} Q_2(\Lambda_g, \psi_g; \boldsymbol{\theta}^{(k+1/2)}) &= C(\boldsymbol{\theta}_1^{(k+1)}) + \frac{1}{2} n_g^{(k+1)} \ln |\psi_g^{-1} \mathbf{I}_p| - \frac{1}{2} n_g^{(k+1)} \psi_g^{-1} \text{tr} \{ \mathbf{S}_g^{(k+1)} \} \\ &\quad + n_g^{(k+1)} \psi_g^{-1} \text{tr} \{ \gamma_g^{(k)} \mathbf{S}_g^{(k+1)} \Lambda_g \} - \frac{1}{2} n_g^{(k+1)} \psi_g^{-1} \text{tr} \{ \Lambda_g \boldsymbol{\Theta}_g^{(k)} \Lambda_g' \}, \end{aligned}$$

yielding

$$\frac{\partial Q_2}{\partial \psi_g^{-1}} = \frac{1}{2} n_g^{(k+1)} \left[p \psi_g - \text{tr} \{ \mathbf{S}_g^{(k+1)} \} + 2 \text{tr} \{ \gamma_g^{(k)} \mathbf{S}_g^{(k+1)} \Lambda_g \} - \text{tr} \{ \Lambda_g \boldsymbol{\Theta}_g^{(k)} \Lambda_g' \} \right].$$

Then the estimated ψ_g is attained for $\hat{\psi}_g$, satisfying

$$\frac{\partial Q_2}{\partial \psi_g^{-1}} = 0 \quad \Rightarrow \quad p \psi_g - \text{tr} \{ \mathbf{S}_g^{(k+1)} \} + 2 \text{tr} \{ \gamma_g^{(k)} \mathbf{S}_g^{(k+1)} \Lambda_g \} - \text{tr} \{ \Lambda_g \boldsymbol{\Theta}_g^{(k)} \Lambda_g' \} = 0.$$

Thus, according to (13), for $\Lambda_g = \hat{\Lambda}_g = \mathbf{S}_g^{(k+1)} \gamma_g^{\prime(k)} \boldsymbol{\Theta}_g^{-1}$ we get $\text{tr} \{ \Lambda_g \boldsymbol{\Theta}_g^{(k)} \Lambda_g' \} = \text{tr} \{ \gamma_g^{(k)} \mathbf{S}_g^{(k+1)} \Lambda_g \}$ and, finally,

$$\hat{\psi}_g = \frac{1}{p} \text{tr} \{ \mathbf{S}_g^{(k+1)} - \hat{\Lambda}_g \gamma_g^{(k)} \mathbf{S}_g^{(k+1)} \}.$$

Thus,

$$\begin{aligned}\psi_g^+ &= \frac{1}{p} \text{tr} \left\{ \mathbf{S}_g^{(k+1)} - \mathbf{\Lambda}_g \gamma_g^+ \mathbf{S}_g^{(k+1)} \right\} \\ \gamma_g^+ &= \mathbf{\Lambda}_g' \left(\mathbf{\Lambda}_g \mathbf{\Lambda}_g' + \psi_g^+ \mathbf{I}_p \right)^{-1},\end{aligned}\tag{B.1}$$

with $\mathbf{\Theta}_g^+$ computed according to (15).

B.2.2 Equal error variance: $\mathbf{\Psi}_g = \mathbf{\Psi}$

In this case, from Equation (10), we have

$$\begin{aligned}Q_2 \left(\mathbf{\Lambda}_g, \mathbf{\Psi}; \boldsymbol{\theta}^{(k+1/2)} \right) &= C(\boldsymbol{\theta}_1^{(k+1)}) - \frac{1}{2} n_g^{(k+1)} \ln |\mathbf{\Psi}| - \frac{1}{2} n_g^{(k+1)} \text{tr} \left\{ \mathbf{S}_g^{(k+1)} \mathbf{\Psi}^{-1} \right\} \\ &\quad + n_g^{(k+1)} \text{tr} \left\{ \mathbf{\Lambda}_g \gamma_g^{(k)} \mathbf{S}_g^{(k+1)} \mathbf{\Psi}^{-1} \right\} - \frac{1}{2} n_g^{(k+1)} \text{tr} \left\{ \mathbf{\Lambda}_g' \mathbf{\Psi}^{-1} \mathbf{\Lambda}_g \mathbf{\Theta}_g^{(k)} \right\},\end{aligned}$$

yielding

$$\frac{\partial Q_2 \left(\mathbf{\Lambda}_g, \mathbf{\Psi}; \boldsymbol{\theta}^{(k+1/2)} \right)}{\partial \mathbf{\Psi}^{-1}} = \frac{1}{2} n_g^{(k+1)} \mathbf{\Psi} - \frac{1}{2} n_g^{(k+1)} \mathbf{S}_g^{(k+1)} + n_g^{(k+1)} \mathbf{S}_g'^{(k+1)} \gamma_g'^{(k)} \mathbf{\Lambda}_g' - \frac{1}{2} n_g^{(k+1)} \mathbf{\Lambda}_g \mathbf{\Theta}_g^{(k)} \mathbf{\Lambda}_g'.$$

Then the estimated $\hat{\mathbf{\Psi}}$ is obtained by satisfying

$$\sum_{g=1}^G \frac{\partial Q_2 \left(\mathbf{\Lambda}_g, \mathbf{\Psi}; \boldsymbol{\theta}^{(k+1/2)} \right)}{\partial \mathbf{\Psi}^{-1}} = \mathbf{0},$$

that is

$$\frac{n}{2} \mathbf{\Psi} - \frac{1}{2} \sum_{g=1}^G n_g^{(k+1)} \mathbf{S}_g^{(k+1)} + \sum_{g=1}^G n_g^{(k+1)} \mathbf{S}_g'^{(k+1)} \gamma_g'^{(k)} \mathbf{\Lambda}_g' - \frac{1}{2} \sum_{g=1}^G n_g^{(k+1)} \mathbf{\Lambda}_g \mathbf{\Theta}_g^{(k)} \mathbf{\Lambda}_g' = \mathbf{0},$$

which can be simplified as

$$\frac{n}{2} \mathbf{\Psi} - \frac{1}{2} \sum_{g=1}^G n_g^{(k+1)} \left[\mathbf{S}_g^{(k+1)} + 2 \mathbf{S}_g'^{(k+1)} \gamma_g'^{(k)} \mathbf{\Lambda}_g' - \mathbf{\Lambda}_g \mathbf{\Theta}_g^{(k)} \mathbf{\Lambda}_g' \right] = \mathbf{0},$$

with $\sum_{g=1}^G n_g^{(k+1)} = n$. Again, according to (13), for $\mathbf{\Lambda}_g = \hat{\mathbf{\Lambda}}_g = \mathbf{S}_g^{(k+1)} \gamma_g'^{(k)} \mathbf{\Theta}_g^{-1}$

we get $\hat{\mathbf{\Lambda}}_g \mathbf{\Theta}_g^{(k)} \hat{\mathbf{\Lambda}}_g' = \hat{\mathbf{\Lambda}}_g \gamma_g^{(k)} \mathbf{S}_g^{(k+1)}$ and, afterwards,

$$\hat{\mathbf{\Psi}} = \sum_{g=1}^G \frac{n_g}{n} \text{diag} \left\{ \mathbf{S}_g^{(k+1)} - \hat{\mathbf{\Lambda}}_g \gamma_g'^{(k)} \mathbf{S}_g^{(k+1)} \right\} = \sum_{g=1}^G \pi_g^{(k+1)} \text{diag} \left\{ \mathbf{S}_g^{(k+1)} - \hat{\mathbf{\Lambda}}_g \gamma_g^{(k)} \mathbf{S}_g^{(k+1)} \right\}.\tag{B.2}$$

Thus,

$$\Psi^+ = \sum_{g=1}^G \pi_g^{(k+1)} \text{diag} \left\{ \mathbf{S}_g^{(k+1)} - \mathbf{\Lambda}_g^+ \boldsymbol{\gamma}_g \mathbf{S}_g^{(k+1)} \right\}, \quad (\text{B.3})$$

$$\boldsymbol{\gamma}_g^+ = \mathbf{\Lambda}_g' \left(\mathbf{\Lambda}_g^+ \mathbf{\Lambda}_g'^+ + \Psi^+ \right)^{-1}$$

where Θ_g^+ is computed according to (15).

B.2.3 Equal loading matrices: $\mathbf{\Lambda}_g = \mathbf{\Lambda}$

In this case, Equation (10) can be written as

$$\begin{aligned} Q_2 \left(\mathbf{\Lambda}, \Psi_g; \boldsymbol{\theta}^{(k+1/2)} \right) = & \text{C}(\boldsymbol{\theta}_1^{(k+1)}) + \frac{1}{2} n_g^{(k+1)} \ln |\Psi_g^{-1}| - \frac{1}{2} n_g^{(k+1)} \text{tr} \left\{ \mathbf{S}_g^{(k+1)} \Psi_g^{-1} \right\} \\ & + n_g^{(k+1)} \text{tr} \left\{ \mathbf{\Lambda} \boldsymbol{\gamma}_g^{(k)} \mathbf{S}_g^{(k+1)} \Psi_g^{-1} \right\} - \frac{1}{2} n_g^{(k+1)} \text{tr} \left\{ \mathbf{\Lambda}' \Psi_g^{-1} \mathbf{\Lambda} \Theta_g^{(k)} \right\}, \end{aligned}$$

yielding

$$\frac{\partial Q_2 \left(\mathbf{\Lambda}, \Psi_g; \boldsymbol{\theta}^{(k+1/2)} \right)}{\partial \mathbf{\Lambda}} = n_g^{(k+1)} \Psi_g^{-1} \mathbf{S}_g^{(k+1)} \boldsymbol{\gamma}_g'^{(k)} - n_g^{(k+1)} \Psi_g^{-1} \mathbf{\Lambda} \Theta_g^{(k)} = \mathbf{0}.$$

Then the estimated $\hat{\mathbf{\Lambda}}$ is obtained by solving

$$\sum_{g=1}^G \frac{\partial Q_2 \left(\mathbf{\Lambda}, \Psi_g; \boldsymbol{\theta}^{(k+1/2)} \right)}{\partial \mathbf{\Lambda}} = \sum_{g=1}^G n_g^{(k+1)} \Psi_g^{-1} \left[\mathbf{S}_g^{(k+1)} \boldsymbol{\gamma}_g'^{(k)} - \mathbf{\Lambda} \Theta_g^{(k)} \right] = \mathbf{0}, \quad (\text{B.4})$$

with $\boldsymbol{\gamma}_g^{(k)} = \mathbf{\Lambda}'^{(k)} \left(\mathbf{\Lambda}^{(k)} \mathbf{\Lambda}'^{(k)} + \Psi_g^{(k)} \right)^{-1}$. In this case, the loading matrix cannot be solved directly and must be solved in a row-by-row manner as suggested by McNicholas and Murphy (2008). Therefore,

$$\lambda_i^+ = \mathbf{r}_i \left(\sum_{g=1}^G \frac{n_g}{\psi_{g(i)}} \Theta_g \right)^{-1} \quad (\text{B.5})$$

$$\boldsymbol{\gamma}_g^+ = \mathbf{\Lambda}' \left(\mathbf{\Lambda}^+ \mathbf{\Lambda}'^+ + \Psi_g^+ \right)^{-1} \quad (\text{B.6})$$

$$\Theta_g^+ = \mathbf{I}_q - \boldsymbol{\gamma}_g^+ \mathbf{\Lambda}^+ + \boldsymbol{\gamma}_g^+ \mathbf{S}_g^{(k+1)} \boldsymbol{\gamma}_g'^+, \quad (\text{B.7})$$

where λ_i^+ is the i th row of the matrix $\mathbf{\Lambda}^+$, $\psi_{g(i)}$ is the i th diagonal element of Ψ_g , and \mathbf{r}_i represents the i th row of the matrix $\sum_{g=1}^G n_g^{(k+1)} \left(\Psi_g' \right)^{-1} \mathbf{S}_g^{(k+1)}$.

B.2.4 Further details

A further schematization is here given without considering the constraint on the Y variable. Thus, with reference to the model identifier, we will only refer

to the last three letters.

Models ended by UUU: no constraint is assumed.

Models ended by UUC: $\Psi_g = \psi_g \mathbf{I}_p$, where the parameter ψ_g is updated according to (B.1).

Models ended by UCU: $\Psi_g = \Psi$, where the matrix Ψ is updated according to (B.3).

Models ended by UCC: $\Psi_g = \psi \mathbf{I}_p$. By combining (B.1) and (B.3) we obtain

$$\hat{\psi} = \frac{1}{p} \sum_{g=1}^G \frac{n_g^{(k+1)}}{n} \text{tr} \left\{ \mathbf{S}_g^{(k+1)} - \hat{\Lambda}_g \gamma_g^{(k)} \mathbf{S}_g^{(k+1)} \right\} = \frac{1}{p} \sum_{g=1}^G \hat{\pi}_g^{(k+1)} \text{tr} \left\{ \mathbf{S}_g^{(k+1)} - \hat{\Lambda}_g \gamma_g^{(k)} \mathbf{S}_g^{(k+1)} \right\}. \quad (\text{B.8})$$

Thus,

$$\begin{aligned} \psi^+ &= \frac{1}{p} \sum_{g=1}^G \pi_g^{(k+1)} \text{tr} \left\{ \mathbf{S}_g^{(k+1)} - \Lambda_g^+ \gamma_g \mathbf{S}_g^{(k+1)} \right\} \\ \gamma_g^+ &= \Lambda_g'^+ \left(\Lambda_g^+ \Lambda_g'^+ + \psi^+ \mathbf{I}_p \right)^{-1}, \end{aligned}$$

with Θ_g^+ computed according to (15).

Models ended by CUU: $\Lambda_g = \Lambda$, where the matrix Λ is updated according to (B.5). In this case, Ψ_g is estimated directly from (12) and thus $\Psi_g^+ = \text{diag} \left\{ \mathbf{S}_g^{(k+1)} - 2\Lambda^+ \gamma_g \mathbf{S}_g^{(k+1)} + \Lambda^+ \Theta_g \Lambda'^+ \right\}$, with γ_g^+ and Θ_g^+ computed according to (B.6) and (B.7), respectively.

Models ended by CUC: $\Lambda_g = \Lambda$ and $\Psi_g = \psi_g \mathbf{I}_p$. In this case, equation (B.4), for $\Psi_g = \psi_g \mathbf{I}_p$, yields

$$\sum_{g=1}^G \frac{\partial Q_2 \left(\Lambda, \psi_g; \boldsymbol{\theta}^{(k+1/2)} \right)}{\partial \Lambda} = \sum_{g=1}^G n_g^{(k+1)} \psi_g^{-1} \mathbf{S}_g^{(k+1)} \gamma_g'^{(k)} - \sum_{g=1}^G n_g^{(k+1)} \psi_g^{-1} \Theta_g^{(k)} = \mathbf{0},$$

and afterwards

$$\hat{\Lambda} = \left(\sum_{g=1}^G \frac{n_g^{(k+1)}}{\psi_g^{-1}} \mathbf{S}_g^{(k+1)} \gamma_g'^{(k)} \right) \left(\sum_{g=1}^G \frac{n_g^{(k+1)}}{\psi_g^{-1}} \Lambda \right)^{-1},$$

with $\gamma_g^{(k)} = \Lambda'^{(k)} \left(\Lambda^{(k)} \Lambda'^{(k)} + \psi_g^{(k)} \mathbf{I}_p \right)^{-1}$. Moreover, from

$$\begin{aligned} \frac{\partial Q_2 \left(\Lambda, \psi_g; \boldsymbol{\theta}^{(k+1/2)} \right)}{\partial \psi_g^{-1}} &= \frac{p}{2} \psi_g - \frac{n_g^{(k+1)}}{2} \left[\text{tr} \left\{ \mathbf{S}_g^{(k+1)} \right\} - 2 \text{tr} \left\{ \mathbf{S}_g'^{(k+1)} \gamma_g'^{(k)} \Lambda' \right\} + \text{tr} \left\{ \Lambda \boldsymbol{\Theta}_g^{(k+1)} \Lambda' \right\} \right] \\ &= 0 \end{aligned}$$

we get $\hat{\psi}_g = (1/p) \text{tr} \left\{ \mathbf{S}_g^{(k+1)} - 2 \hat{\Lambda} \gamma_g'^{(k)} \mathbf{S}_g + \hat{\Lambda} \boldsymbol{\Theta}_g \hat{\Lambda}' \right\}$. Thus,

$$\begin{aligned} \Lambda^+ &= \left(\sum_{g=1}^G \frac{n_g^{(k+1)}}{\psi_g^{-1}} \mathbf{S}_g^{(k+1)} \gamma_g' \right) \left(\sum_{g=1}^G \frac{n_g^{(k+1)}}{\psi_g^{-1}} \Lambda \right)^{-1} \\ \psi_g^+ &= \frac{1}{p} \text{tr} \left\{ \mathbf{S}_g^{(k+1)} - 2 \Lambda^+ \gamma_g' \mathbf{S}_g + \Lambda^+ \boldsymbol{\Theta} \Lambda'^+ \right\} \\ \gamma_g^+ &= \Lambda'^+ \left(\Lambda^+ \Lambda'^+ + \psi_g^+ \mathbf{I}_p \right)^{-1}. \end{aligned}$$

with $\boldsymbol{\Theta}_g^+$ computed according to (B.7).

Models ended by CCU: $\Lambda_g = \Lambda$ and $\Psi_g = \Psi$, so that $\gamma^{(k)} = \Lambda'^{(k)} \left(\Lambda^{(k)} \Lambda'^{(k)} + \Psi^{(k)} \right)^{-1}$.

Setting $\Psi_g = \Psi$ in (B.4), we get

$$\begin{aligned} \sum_{g=1}^G \frac{\partial Q_2 \left(\Lambda, \Psi; \boldsymbol{\theta}^{(k+1/2)} \right)}{\partial \Lambda} &= \sum_{g=1}^G n_g^{(k+1)} \Psi^{-1} \left[\mathbf{S}_g^{(k+1)} \gamma'^{(k)} - \Lambda \boldsymbol{\Theta}_g^{(k)} \right] \\ &= \Psi^{-1} \left[\gamma'^{(k)} \sum_{g=1}^G n_g^{(k+1)} \mathbf{S}_g^{(k+1)} - \Lambda \sum_{g=1}^G n_g^{(k+1)} \boldsymbol{\Theta}_g^{(k)} \right] \\ &= \Psi^{-1} \left[\gamma'^{(k)} \mathbf{S}^{(k+1)} - \Lambda \boldsymbol{\Theta}^{(k)} \right] = \mathbf{0}, \end{aligned}$$

where

$$\begin{aligned} \mathbf{S}^{(k+1)} &= \sum_{g=1}^G \pi_g^{(k+1)} \mathbf{S}_g^{(k+1)} \\ \boldsymbol{\Theta}^{(k)} &= \sum_{g=1}^G \pi_g^{(k+1)} \boldsymbol{\Theta}_g^{(k)} = \mathbf{I}_q - \gamma^{(k)} \Lambda^{(k)} + \gamma^{(k)} \mathbf{S}^{(k+1)} \gamma'^{(k)}. \end{aligned}$$

Thus,

$$\hat{\Lambda} = \mathbf{S}^{(k+1)} \gamma'^{(k)} \left(\boldsymbol{\Theta}^{(k)} \right)^{-1}. \quad (\text{B.9})$$

Moreover, setting $\Lambda_g = \Lambda$ in (B.2), we get $\hat{\Psi} = \text{diag} \left\{ \mathbf{S}^{(k+1)} - \hat{\Lambda} \gamma^{(k)} \mathbf{S}^{(k+1)} \right\}$. Hence,

$$\begin{aligned} \Lambda^+ &= \mathbf{S}^{(k+1)} \gamma' \boldsymbol{\Theta}^{-1} \\ \Psi^+ &= \text{diag} \left\{ \mathbf{S}^{(k+1)} - \Lambda^+ \gamma \mathbf{S}^{(k+1)} \right\} \\ \gamma_g^+ &= \Lambda'^+ \left(\Lambda^+ \Lambda'^+ + \Psi^+ \right)^{-1}, \end{aligned} \quad (\text{B.10})$$

with Θ_g^+ computed according to (B.7).

Models ended by CCC: $\Lambda_g = \Lambda$ and $\Psi_g = \psi \mathbf{I}_p$, so that $\gamma^{(k)} = \Lambda'^{(k)} \left(\Lambda^{(k)} \Lambda'^{(k)} + \psi^{(k)} \right)^{-1}$.

Here, the estimated loading matrix is again (B.9), while the isotropic term obtained from (B.8) for $\Lambda_g = \Lambda$ is $\hat{\psi} = (1/p) \text{tr} \left\{ \mathbf{S}^{(k+1)} - \hat{\Lambda} \gamma^{(k)} \mathbf{S}^{(k+1)} \right\}$,

with $\gamma_g^{(k)} = \Lambda_g'^{(k)} \left(\Lambda_g^{(k)} \Lambda_g'^{(k)} + \psi^{(k)} \mathbf{I}_p \right)^{-1}$. Hence,

$$\begin{aligned} \psi^+ &= \frac{1}{p} \text{tr} \left\{ \mathbf{S}^{(k+1)} - \Lambda^+ \gamma \mathbf{S}^{(k+1)} \right\} \\ \gamma^+ &= \Lambda'^+ \left(\Lambda^+ \Lambda'^+ + \psi^+ \mathbf{I}_p \right)^{-1}, \end{aligned}$$

with Λ^+ and Θ_g^+ computed according to (B.10) and (B.7), respectively.

C True and estimated covariance matrices of Section 6.1

Because the loading matrices are not unique, for the simulated data of Examples 1 and 2 we limit the attention to a comparison, for each $g = 1, \dots, G$, of true and estimated covariance matrices.

C.1 Example 6.1.1

$$\Sigma_1 = \begin{bmatrix} 103.36 & 103.07 & 101.37 & 79.41 & 105.66 \\ 103.08 & 119.39 & 110.23 & 85.97 & 115.47 \\ 101.37 & 110.23 & 129.77 & 106.08 & 118.50 \\ 79.41 & 85.97 & 106.08 & 101.46 & 95.21 \\ 105.66 & 115.47 & 118.50 & 95.21 & 121.63 \end{bmatrix} \quad \hat{\Sigma}_1 = \begin{bmatrix} 107.59 & 114.55 & 110.42 & 87.29 & 114.43 \\ 114.55 & 139.40 & 127.06 & 100.09 & 132.06 \\ 110.42 & 127.06 & 146.31 & 122.92 & 134.12 \\ 87.29 & 100.09 & 122.92 & 117.97 & 110.09 \\ 114.43 & 132.06 & 134.12 & 110.09 & 135.66 \end{bmatrix},$$

and

$$\Sigma_2 = \begin{bmatrix} 34.25 & 15.16 & 17.81 & 22.39 & 14.62 \\ 15.16 & 17.01 & 11.42 & 13.98 & 8.95 \\ 17.81 & 11.42 & 17.62 & 16.12 & 10.45 \\ 22.39 & 13.98 & 16.12 & 28.11 & 13.11 \\ 14.62 & 8.95 & 10.45 & 13.11 & 10.19 \end{bmatrix} \quad \hat{\Sigma}_2 = \begin{bmatrix} 22.16 & 7.44 & 13.71 & 12.89 & 10.12 \\ 7.44 & 11.25 & 7.59 & 8.05 & 5.48 \\ 13.71 & 7.59 & 18.83 & 13.53 & 10.13 \\ 12.89 & 8.05 & 13.53 & 22.00 & 9.41 \\ 10.12 & 5.48 & 10.13 & 9.41 & 8.63 \end{bmatrix}.$$

C.2 Example 6.1.2

$$\mathbf{\Sigma}_1 = \begin{bmatrix} 10.41 & 3.61 & 4.07 & 4.48 & 5.71 \\ 3.61 & 7.83 & 2.88 & 3.18 & 4.03 \\ 4.07 & 2.88 & 8.67 & 3.81 & 4.64 \\ 4.48 & 3.18 & 3.81 & 9.61 & 5.17 \\ 5.71 & 4.04 & 4.64 & 5.17 & 11.73 \end{bmatrix} \quad \hat{\mathbf{\Sigma}}_1 = \begin{bmatrix} 8.86 & 3.89 & 5.06 & 3.84 & 5.72 \\ 3.89 & 7.23 & 3.59 & 1.79 & 4.04 \\ 5.06 & 3.59 & 8.44 & 3.85 & 5.50 \\ 3.84 & 1.79 & 3.85 & 7.74 & 4.38 \\ 5.72 & 4.04 & 5.50 & 4.38 & 9.81 \end{bmatrix},$$

$$\mathbf{\Sigma}_2 = \begin{bmatrix} 103.36 & 103.07 & 101.37 & 79.41 & 105.66 \\ 103.08 & 122.1 & 110.23 & 85.97 & 115.47 \\ 101.37 & 110.23 & 134.33 & 106.08 & 118.50 \\ 79.41 & 85.97 & 106.08 & 102.73 & 95.21 \\ 105.66 & 115.47 & 118.50 & 95.21 & 129.21 \end{bmatrix} \quad \hat{\mathbf{\Sigma}}_2 = \begin{bmatrix} 106.17 & 100.46 & 93.18 & 73.81 & 105.01 \\ 100.46 & 113.71 & 92.97 & 72.22 & 107.88 \\ 93.18 & 92.97 & 108.25 & 83.08 & 102.36 \\ 73.81 & 72.22 & 83.08 & 80.09 & 81.85 \\ 105.01 & 107.88 & 102.36 & 81.85 & 122.59 \end{bmatrix},$$

and

$$\mathbf{\Sigma}_3 = \begin{bmatrix} 25.19 & 15.16 & 17.81 & 22.39 & 14.62 \\ 15.16 & 10.67 & 11.42 & 13.98 & 8.95 \\ 17.81 & 11.42 & 13.12 & 16.12 & 10.45 \\ 22.39 & 13.98 & 16.12 & 20.31 & 13.11 \\ 14.62 & 8.95 & 10.45 & 13.11 & 8.70 \end{bmatrix} \quad \hat{\mathbf{\Sigma}}_3 = \begin{bmatrix} 32.47 & 19.91 & 23.06 & 28.78 & 18.80 \\ 19.91 & 14.10 & 14.96 & 18.25 & 11.66 \\ 23.06 & 14.96 & 16.95 & 20.77 & 13.45 \\ 28.78 & 18.25 & 20.77 & 25.95 & 16.77 \\ 18.80 & 11.66 & 13.45 & 16.77 & 11.10 \end{bmatrix}.$$

References

- Airolidi, J. and R. Hoffmann (1984). Age variation in voles (*Microtus californicus*, *M. ochrogaster*) and its significance for systematic studies. Occasional papers of the Museum of Natural History 111, University of Kansas, Lawrence, KS.
- Aitken, A. C. (1926). On Bernoulli's numerical solution of algebraic equations. In *Proceedings of the Royal Society of Edinburgh*, Volume 46, pp. 289–305.
- Banfield, J. D. and A. E. Raftery (1993). Model-based Gaussian and non-Gaussian clustering. *Biometrics* 49(3), 803–821.
- Bartlett, M. (1953). Factor analysis in psychology as a statistician sees it. In *Uppsala symposium on psychological factor analysis*, Number 3 in Nordisk

- Psykologi's Monograph Series, Uppsala, Sweden, pp. 23–34. Almquist and Wiksell Uppsala.
- Biernacki, C., G. Celeux, and G. Govaert (2000). Assessing a mixture model for clustering with the integrated completed likelihood. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 22(7), 719–725.
- Böhning, D., E. Dietz, R. Schaub, P. Schlattmann, and B. Lindsay (1994). The distribution of the likelihood ratio for mixtures of densities from the one-parameter exponential family. *Annals of the Institute of Statistical Mathematics* 46(2), 373–388.
- Carvalho, C., J. Chang, J. Lucas, J. Nevins, Q. Wang, and M. West (2008). High-dimensional sparse factor modeling: applications in gene expression genomics. *Journal of the American Statistical Association* 103(484), 1438–1456.
- Celeux, G. and G. Govaert (1995). Gaussian parsimonious clustering models. *Pattern Recognition* 28(5), 781–793.
- Dempster, A., N. Laird, and D. Rubin (1977). Maximum likelihood from incomplete data via the EM algorithm. *Journal of the Royal Statistical Society. Series B (Methodological)* 39(1), 1–38.
- Everitt, B. and D. J. Hand (1981). *Finite mixture distributions*. Chapman & Hall.
- Flury, B. (1997). *A first course in multivariate statistics*. New York: Springer.
- Fraley, C. and A. E. Raftery (2002). Model-based clustering, discriminant analysis, and density estimation. *Journal of the American Statistical Association* 97(458), 611–631.
- Fraley, C., A. E. Raftery, T. B. Murphy, and L. Scrucca (2012). *mclust* version 4 for R: Normal mixture modeling for model-based clustering, classification, and density estimation. Technical report 597, Department of Statistics, University of Washington, Seattle, Washington, USA.
- Frühwirth-Schnatter, S. (2006). *Finite mixture and Markov switching models*. New York: Springer.
- Gershensfeld, N. (1997). Nonlinear inference and cluster-weighted modeling. *Annals of the New York Academy of Sciences* 808(1), 18–24.
- Ghahramani, Z. and G. Hinton (1987). The EM algorithm for factor analyzers. Technical report CRG-TR-96-1, University of Toronto.
- Hosmer Jr., D. (1973). A comparison of iterative maximum likelihood estimates of the parameters of a mixture of two normal distributions under three different types of sample. *Biometrics* 29(4), 761–770.
- Hubert, L. and P. Arabie (1985). Comparing partitions. *Journal of Classification* 2(1), 193–218.
- Ingrassia, S., S. C. Minotti, and A. Punzo (2012). Model-based clustering via linear cluster-weighted models. eprint arXiv: 1206.3974, <http://arxiv.org/abs/1206.3974>.
- Ingrassia, S., S. C. Minotti, and G. Vittadini (2012). Local statistical modeling via the cluster-weighted approach with elliptical distributions. *Journal of Classification* 29(3). 363–401.

- McLachlan, G., D. Peel, and R. Bean (2003). Modelling high-dimensional data by mixtures of factor analyzers. *Computational Statistics & Data Analysis* 41(3), 379–388.
- McLachlan, G. J. and D. Peel (2000). *Finite Mixture Models*. New York: John Wiley & Sons.
- McNicholas, P. (2010). Model-based classification using latent Gaussian mixture models. *Journal of Statistical Planning and Inference* 140(5), 1175–1181.
- McNicholas, P. D., K. R. Jampani, A. F. McDaid, T. B. Murphy, and L. Banks (2011). *pgmm: Parsimonious Gaussian Mixture Models*. R package version 1.0.
- McNicholas, P. D. and T. B. Murphy (2008). Parsimonious Gaussian mixture models. *Statistics and Computing* 18(3), 285–296.
- McNicholas, P. D., T. B. Murphy, A. F. McDaid, and D. Frost (2010). Serial and parallel implementations of model-based clustering via parsimonious Gaussian mixture models. *Computational Statistics & Data Analysis* 54(3), 711–723.
- Meng, X. L. and D. van Dyk (1997). The EM algorithm – an old folk-song sung to a fast new tune. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* 59(3), 511–567.
- Press, W., S. Teukolsky, W. Vetterling, and B. Flannery (1992). *Numerical recipes in C - The Art of Scientific Computation* (2nd ed.). Cambridge: Cambridge University Press.
- Punzo, A. (2012). Flexible mixture modeling with the polynomial Gaussian cluster-weighted model. eprint arXiv: 1207.0939, <http://arxiv.org/abs/1207.0939>.
- Rand, W. (1971). Objective criteria for the evaluation of clustering methods. *Journal of the American Statistical Association* 66(336), 846–850.
- Sakamoto, Y., M. Ishiguro, and G. Kitagawa (1983). *Akaike information criterion statistics*. Boston: Reidel.
- Schwarz, G. (1978). Estimating the dimension of a model. *The Annals of Statistics* 6(2), 461–464.
- R Development Core Team (2012). *R: A Language and Environment for Statistical Computing*. Vienna, Austria: R Foundation for Statistical Computing.
- Spearman, C. (1904). The proof and measurement of association between two things. *The American Journal of Psychology* 15(1), 72–101.
- Titterton, D. M., A. F. M. Smith, and U. E. Makov (1985). *Statistical Analysis of Finite Mixture Distributions*. New York: John Wiley & Sons.
- Wang, Q., C. Carvalho, J. Lucas, and M. West (2007). BFRM: Bayesian factor regression modelling. *Bulletin of the International Society for Bayesian Analysis* 14(2), 4–5.
- West, M. (2003). Bayesian factor regression models in the “large p , small n ” paradigm. In J. Bernardo, M. Bayarri, J. Berger, A. Dawid, D. Heckerman, A. Smith, and M. West (Eds.), *Bayesian statistics*, Volume 7, Oxford, pp. 723–732. Oxford University Press.

Woodbury, M. A. (1950). *Inverting modified matrices*. Statistical Research Group, Memo. Rep. no. 42. Princeton, New Jersey: Princeton University.